Official Transcript of Proceedings

NUCLEAR REGULATORY COMMISSION

Title: Advisory Committee on Reactor Safeguards
Digital Instrumentation and Control
Systems Subcommittee

Docket Number: (n/a)

Location: Rockville, Maryland

Date: Wednesday, August 19, 2009

Work Order No.: NRC-3021 Pages 1-366

NEAL R. GROSS AND CO., INC.
Court Reporters and Transcribers
1323 Rhode Island Avenue, N.W.
Washington, D.C. 20005
(202) 234-4433

DISCLAIMER

UNITED STATES NUCLEAR REGULATORY COMMISSION'S ADVISORY COMMITTEE ON REACTOR SAFEGUARDS

The contents of this transcript of the proceeding of the United States Nuclear Regulatory Commission Advisory Committee on Reactor Safeguards, as reported herein, is a record of the discussions recorded at the meeting.

This transcript has not been reviewed, corrected, and edited, and it may contain inaccuracies.

UNITED STATES OF AMERICA

NUCLEAR REGULATORY COMMISSION

+ + + + +

ADVISORY COMMITTEE ON REACTOR SAFEGUARD

(ACRS)

+ + + + +

SUBCOMMITTEE ON DIGITAL INSTRUMENTATION AND

CONTROL SYSTEMS

+ + + + +

WEDNESDAY

AUGUST 19, 2009

+ + + + +

ROCKVILLE, MARYLAND

+ + + + +

The Subcommittee convened at the Nuclear Regulatory Commission, Two White Flint North, Room T-2B3, 11545 Rockville Pike, at 8:30 a.m., Dr. George Apostolakis, Chairman, presiding.

SUBCOMMITTEE MEMBERS:

GEORGE APOSTOLAKIS, Chairman

SAID ABDEL-KHALIK, Member

DENNIS C. BLEY, Member

CHARLES H. BROWN, JR., Member

JOHN D. SIEBER, Member

JOHN W. STETKAR, Member

CONSULTANT TO THE SUBCOMMITTEE:

MYRON HECHT

NRC STAFF PRESENT:

CHRISTINA ANTONESCU, Cognizant Staff Engineer and Designated Federal Official

PAT HILAND

DAN SANTOS

MICHAEL WATERMAN

ALAN KURITZKY

ALSO PRESENT:

ROB AUSTIN

RAY TOROK

THUY NGUYEN

BRUCE GEDDES

DAVE BLANCHARD

C-O-N-T-E-N-T-S

Opening Remarks	4
George Apostolakis, Chair	
Introductory Remarks	5
Patrick Hiland	
Director, Division of Engineering	
Nuclear Reactor Regulation	
EPRI Presentation on Operating	6
Experience Report, DAS Report, and	
Failure Modes Research	
Rob Austin	6
Ray Torok	16, 216, 259, 354
Bruce Geddes	83, 177, 217
Dave Blanchard	162, 272
Thuy Nguyen	231

P-R-O-C-E-E-D-I-N-G-S

8:31 a.m.

CHAIR APOSTOLAKIS: The meeting will now come to order.

This is a meeting of the Digital Instrumentation and Control Systems Subcommittee of the Advisory Committee on Reactor Safeguards. I am George Apostolakis, Chairman of the Subcommittee.

ACRS members in attendance are Dennis Bley, John Stetkar, Jack Sieber, and Charles Brown. Myron Hecht is also attending as a consultant for the Subcommittee.

Christina Antonescu of the ACRS staff is a Designated Federal Official for this meeting.

The purpose of this meeting is to discuss the Draft Interim Staff Guidance No. 6 on licensing process and Draft ISG No. 7 on fuel facilities. We will also discus the digital I&C research plan for fiscal year 2010 to 2014 and EPRI's reports on operating experience and diverse actuation systems risks and benefits.

The Subcommittee will gather information, analyze relevant issues and facts, and formulate proposed positions and actions as appropriate for deliberation by the full Committee.

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 July 21st, 2009.

We have received no written comments or requests for time to make oral statements from members of the public regarding today's meeting.

We have several people on the bridge phone line listening to the discussions. To preclude interruption of the meeting, the phone line will be placed on listen-in mode during the discussions, presentations, and Committee deliberations.

A transcript of the meeting is being kept and will be made available as stated in The Federal Register notice. Therefore, we request the participants in this meeting use the microphones located throughout the meeting room when addressing the Subcommittee. The participants should first identify themselves and speak with sufficient clarity and volume so that they may be readily heard.

We will now proceed with the meeting. I call upon Mr. Patrick Hiland, Director, Division of Engineering, Nuclear Reactor Regulation, to provide some introductory remarks.

MR. HILAND: Thank you, sir.

Good morning. As I see, you have a very full agenda over the next three days. Normally, Mr. Jack Grobe would give some introductory remarks to this Subcommittee. However, he is busy with a new assignment that goes through the end of the month in the Office of Nuclear Materials Safety and Safeguards.

What I would like to do first is thank the Electric Power Research Institute for their efforts. Just to let you know, they did come in and meet with the staff earlier in the month, the first week of August. They reviewed with us their draft reports in detail.

We had a very healthy meeting, very good discussion. I believe our reviews of those reports are complete. Initially, we provided some comments earlier this week to the Committee. We have not had a chance to sit down with EPRI on those final comments and discuss some of the questions they may have on our conclusions. So we look forward to doing that with EPRI.

With that, I would like to turn over the presentation to Mr. Rob Austin from EPRI.

MR. AUSTIN: Thanks.

Good morning. I am Robert Austin, INC Program Manager for the Electric Power Research

Institute.

I would like to thank you for the opportunity to present the results of some of EPRI's research in the digital instrumentation and control systems. We have been doing significant research on the subject and we have some very interesting work that we would like to share with you today.

Our purpose in speaking to you is to gain your insights and reaction to it and use this feedback to further inform additional research. We also would like to present our research to industry as ready for application in the plants, but are interested in your reaction to it prior to this step.

I would like to begin with a hypothesis. Digital instrumentation and control systems are more reliable than analog circuit-based systems currently in many of the U.S. commercial and nuclear plants. The designs of digital instrumentation and control systems, combined with rare, yet potential failures, do not introduce consequences any more severe than the consequences of failures of the existing analog system. Therefore, digital instrumentation and control systems are safer.

Looking outside our own industry, we can, of course, observe this hypothesis may have some

merit. Practically every high-technology industry -- aerospace, pharmaceutical, processed chemical, aviation -- has moved to digital INC as a standard. Some international nuclear plants, as you know, have adopted large-scale digital INC as well, and there are some installed systems within the U.S. commercial fleet that have been a success.

There have been some events, which we will discuss later as part of our first topic, but it is important to note that, after installing a digital INC system, very few, perhaps no one, has decided to replace it with an analog-based control system.

Cursory examination of the evidence shows we must be on to something here. So, when we have a hypothesis, of course, you must provide evidence, either analytical or experiential-based, in order to validate your hypothesis and have it become a working theory.

So EPRI is going to present some of the evidence we have developed for your review today. Our evidence has been, and will continue to be, subject to the scrutiny of peer review, both internal and external. This review is part of the normal process of scientific discovery, and we welcome this opportunity to refine our methods and conclusions as

results. However, we believe the final result is that our hypothesis, methodology, and results will be found to be sound.

Today we have three topics related to digital INC that we want to present. First, we will present our analysis of 322 digital system events in the U.S. This analysis shows that, despite the inherent complexity of control systems, software common-cause failure is not prevalent. Where it is found, it is typically the result of errors in the application level of the code. This result is simply another way of saying that design errors can and do happen. But the addition of additional complexity may not avert this problem.

We expect this database to be very useful in informing industry regulatory guidance and future research on failures in digital systems best practices, and we do have plans to expand it to international events in 2010 and 2011.

As we discussed operating experience, we found that the definitions of software, defect, failure, among other terms, are essential. As attributed to Voltaire, "If you wish to converse with me, define your terms." The definitions we used are from previous guidance and were the subject of a lot

of internal debate. We are going to dwell upon them at some length today to make sure that there's clear understanding.

We found it very difficult to have meaningful discussions on operating experience without having a common terminology. So I hope you are able to see the value of our definitions and endorse them for use in future projects of a similar nature or at least endorse the need for common, well-understood, and agreed-to terminology.

Secondly, we will show you some preliminary results of research and the ways to avoid design errors, including common-cause failure precursors, through what we call defensive measures. These methods offer ways to address design errors without adding complexity, which could serve as a source for more design errors. This research will be completed next year, and we would welcome NRC participation in developing a comprehensive and useful list of defensive measures for use by INC system designs and regulators.

Finally, we will show that INC for some applications is a minimal contribution to plant risk, in addition to any diverse systems, that is to say, additional complexity must be done with great caution

to avoid negating the safety benefits.

From a risk perspective, the plants already have significant defense in-depth to cope with design-basis events. Any additional defense in-depth and diversity that may be needed to address potential software common-cause failures should consider this existing defense in-depth as well as the frequency and consequences of the events in question.

A much broader result here is that traditional PRA methods can provide significant risk insights into INC architecture decisions, such as whether or not to have an automated diverse actuation systems, without having all of the details of the INC system actually in the PRA model.

We believe that these methods do not conflict with existing NRC policy on these subjects. We hope you will agree and encourage staff and industry to support this use of risk-informed methods for digital INC designs.

The overall conclusion of the research that we will present today is that the system designer can obtain significant insights from the application of operating experience, defensive measures, and risk-informed methods to INC system design. Application of this research will result in an even more robust

system than those the nuclear industry has installed to date, which have proved quite robust.

This research we hope will enable broader adoption of digital technology for INC and nuclear plants in the U.S., as has already occurred in our peer high-technology industries. As mentioned, we are anxious to hear your comments and reaction to this research. Out intent is to improve our research, address any gaps, and allow the industry to use our research in their dealings with the NRC staff. Our hope is that you will agree that our defense of our hypothesis is fundamentally sound and can serve as the basis for future regulatory activities.

I would like to introduce our primary speakers for each topic.

Ray Torok, whom you know, is our Senior Project Manager responsible within EPRI for research related to digital system design, diversity, and risk. He will be providing the overall results and summaries.

For operating experience analysis, Bruce Geddes of Southern Engineering Services, will provide the presentation of our data sources, methods, and results.

Thuy Nguyen, Electricite de France, will

discuss the use of defensive measures to prevent and mitigate against common-cause failures.

Finally, for our research related to the use of risk insights in the digital system design, David Blanchard, of Applied Reliability Engineering, will present.

Again, thank you for the opportunity to present to you today.

CHAIR APOSTOLAKIS: Is there a cooperative agreement with EPRI on this effort? Are you planning to have one?

MR. SANTOS: This is Dan Santos from the NRC Office of Research.

The answer is, yes, we have entered into a Memorandum of Understanding with EPRI back in March of this year. We had several meetings. We are trying to formulate potential collaborative activities in the near future.

CHAIR APOSTOLAKIS: The work we are hearing today is just EPRI?

MR. SANTOS: Correct.

CHAIR APOSTOLAKIS: But maybe in the future, there will be some collaboration?

MR. SANTOS: Right, and the staff is looking at leveraging some of the work that you are

going to hear today.

MR. AUSTIN: Of course, some of the comments on the research today will help inform those efforts under the MOU with NRC Research.

CHAIR APOSTOLAKIS: Sure. Okay.

Charlie?

MEMBER BROWN: Yes, go back to your first page. You made a rather broad, general statement relative to digital INC. It was like your first or second sentence, that it is more reliable, whatever.

MR. AUSTIN: Yes, sir.

MEMBER BROWN: Could you repeat that?

MR. AUSTIN: Digital INC control systems were more reliable than analog circuit-based systems currently in many of the U.S. commercial nuclear plants.

MEMBER BROWN: Okay. I don't disagree with that. If you're worrying, think I am going to sit up here and disagree, I don't. Okay?

But one of the things, based on your subsequent statements, that really has to be brought into that statement to make that reliable from protecting the reactor, it is response, not just the fact that it doesn't drift as much. It is easier to maintain alignments. There's a whole lot of positive

benefits, self-testing capabilities, on and on and on, that the digital INC brings to the game that analog didn't.

One of the things, though, you are not reliable if you don't maintain part of what I call the four pillars of reliable instrumentation for protective reactor plants. That is redundancy, independence, determinate behavior, and then we can work down into the diversity and defense in-depth aspect, and the last, what I call plus-one, is you like nice, simple software. If you don't have simple software, then you start stepping back and you're walking backwards against this reliability issue.

I just wanted to make sure you understood a perspective --

MR. AUSTIN: Yes.

MEMBER BROWN: -- that if you don't have those particular aspects involved in these designs, and you addressed some of that in here when you were talking about COT systems back in the diverse actuating, the diverse system applications, about COT systems. If you don't have those, then you bring fundamental problems into the aspects of how these things were applied.

But I just wanted to give you a

perspective relative to those because I have only been on the Committee now for about a year and a half. As I am learning more and more and more as I see the systems being applied, independence seems to be, I don't want to say compromised, but it is less easy to discern that you have true independence, based on the way some of these platforms are being applied from channel to channel to channel.

So I just wanted to keep that in mind. Other than that, I just had to get my two cents in.

MR. AUSTIN: I would say it is always the caveat more reliable when correctly applied.

MEMBER BROWN: Thank you. Okay.

Thanks, George.

CHAIR APOSTOLAKIS: Okay.

MR. TOROK: Any questions?

CHAIR APOSTOLAKIS: Go ahead. Ask those questions. They'll move on.

(Laughter.)

MR. TOROK: Oh, I thought we were done there for a second.

We don't disagree with anything you said there. Hopefully, you will see that the message that we bring is pretty consistent with that.

MEMBER BROWN: Well, I am going to be

addressing this in one of our later meetings, which you all won't be in. But, I mean, relative to the independence issue of how these systems are applied. I am not saying the Committee agrees with my thoughts. That's me. I have to talk to them about that.

MR. SIEBER: You're okay so far.

MEMBER BROWN: Thank you.

(Laughter.)

MR. TOROK: Let's try to get into this then.

My name is Ray Torok. As Rob said, I am the EPRI Project Manager on these projects we are talking about today.

For starters, I want to say it's really good to be back. We have been talking about coming back to you guys since, I think, April last year.

I appreciate all the time you have given us on the agenda, so that we can get into some detail. To make sure we don't waste your time on that, I just want to make it clear that we have brought our A team, so that we can respond to your questions at whatever level of detail you want to go to.

Those are the guys who Rob listed earlier: Bruce Geddes, Thuy Nguyen, and Dave Blanchard. They will be presenting the materials in their areas a

little later. So you save your hard questions for them, I guess is what I am trying to say.

Now Rob mentioned, why are we here? Well, there are some activities that we have been involved in for some time now. The main idea is, of course, operating experience review. We have given you a report on that. The details are on the slide. We don't need to discuss that.

What we want to do in terms of this operating experience is sort of pick up where we left off last year in April. As I recall, at the time we stopped, you guys were asking a lot of questions about failure modes and what we learned in the operating experience about digital failure modes, and so on. And we ran out of time. About then, you guys said, "Geez, we're just getting to the good part." So we want to try to take up where we left off there.

Now, to do that, it has been a long time since, so we have to review some of the other material. So we've got some review, and then we want to take it from there.

But this whole discussion of failure modes is one we want to expand on. Then that becomes a common thread through the whole presentation.

So what did we learn about failure modes

from the operating experience, and what about digital system behaviors, in light of failure modes? We have expanded it. You see here it says, "Mechanisms, Modes, and Effects". So there is a lot more discussion of those things.

Of course, the next step is, what does that mean in the PRA world? Where do you go with this whole discussion of mechanisms, modes, and effects? So we want to expand that.

Now a lot of this, well, nearly all of it I guess, really is in response to requests for information that came from this panel back in March and April last year, and there was a letter earlier than that, in fact, really stressing the importance of digital behaviors in regard to common-cause failure, defense in-depth, and diversity, and so on. So we are trying to respond to that and to the issues you guys have identified as important, like this whole failure modes thing. All right? So that's where we want to go with this thing.

As Rob pointed out, we would like to gather feedback from you guys, which is going to help us and help aim our future work on this subject. Now I know you guys are shy about sharing your opinions, but I want to encourage you to tell us what you think

as we move forward.

CHAIR APOSTOLAKIS: Ray?

MR. TOROK: Yes?

CHAIR APOSTOLAKIS: The operating experience, could you remind us to what extent digital INC is being used in reactors?

MR. TOROK: Oh, my, I don't know that we have any specific data on how many systems are out there. In many plants, many, many plants, I guess, they've gone to digital upgrades of frontline control systems that have been problematic in the past. That means feedwater control has been a big one, where there have been a lot of gains in terms of reliability, and resultant gains in safety as well because of the implementation of digital feedwater systems that are far more robust than their analog predecessors.

The same thing with digital EHC, the Electro-Hydraulic Control for the turbine. Those are the two good examples on the non-safety side.

Now, on the safety side, there have been some digital implementations of our RPS many years ago now. There were some Eagle-21 systems put in.

More recently, there have been a few because there's been a lot of controversy over the

details of the NRC reviews, and so on.

Let's see, beyond that, there have been a number of, I guess, piece-part digital upgrades to replace problematic components in systems, those kinds of things.

CHAIR APOSTOLAKIS: Jack?

MR. SIEBER: Yes, I concur with your analysis that, as of two years ago, there were 38 systems that were installed of a digital nature. The majority were three-element feedwater control systems, turbine control systems, rod position indication, which is one you didn't mention.

MR. TOROK: Right.

MR. SIEBER: It doesn't have a control function, but it is important from the standpoint of reactor safety.

I examined LERs in INC systems for the last three years. The digital systems have earned their share of the LERs for mal-operation, particularly in three-element feedwater control and turbine control systems.

So I haven't finished my analysis, but I would say it is sort of like the old days in the analog systems. They do fail.

There are a couple of things that I would

point out that I ran across that we may want to pay attention to. Maybe it is not in our scope, but I have talked to a number of digital INC design engineers in industry outside of my role as an ACRS member. It turns out that there are water protective relays that are in timers and things like this that are digital that can be used, and in some cases are used, in applications in nuclear power plants that I think are significant.

They come with their own list of problems, one of which earned me a civil penalty a number of years ago for misoperation of diesel generator start and load circuits, which was one of the early single applications which was difficult to diagnose because it wouldn't occur during normal operation.

So I basically can confirm, from my own independent research, that this is where the applications are. If you want to look at more complex and more to the nuclear safety application, you actually have to look at Europe and Japan, in particular. That is a worthwhile study.

MEMBER BROWN: The LERs that you said they had their own fair share, were those hardware-type component failures or --

MR. SIEBER: Most of them were hardware,

but not all.

MR. TOROK: We will show you our results on that, too.

MEMBER BROWN: No, no, you have some of that in your OE discussion.

MR. TOROK: Okay. That's right. That's right.

MR. SIEBER: Generally, feedwater controls and turbine controls are relatively simple, don't have a lot of external elements to them. But, since they, between the two of them, represent a fair portion of the digital applications in this country, they get their fair share of the LERs.

MR. TOROK: Right, and, of course, those systems, the turbine control and digital -- or the analog feedwater were real targets for improvement because they had lots of single points of failure that were causing problems.

MR. SIEBER: Yes.

MR. TOROK: Now I know of one plant where they installed their first digital feedwater system I want to say around 1990. At the time, they went through and made a list of all the problems they had ever had with the analog system. They literally tried to design all those problems out with the new digital

system, so it couldn't have those failure modes. And they succeeded in nearly every case.

Of course, that system, according to their accounting, that system paid for itself in the first startup by avoiding some trips that they would have had otherwise. So there are some good stories like that.

MR. SIEBER: Yes, they could have accomplished the same thing with an analog system. One of the difficulties I had with the LERs is it would describe the event very well and say the system failed, and the corrective action is we replaced a card, and I sit and scratch my head as to really what went wrong.

(Laughter.)

MR. TOROK: That is a good point. We had that problem, too. We will talk more about that.

Okay. So did we answer the question?

Okay. Now, just as a bit of background, the question comes up, well, what's EPRI's role in the world here? Of course, we work for the utility industry, so we are trying to help them improve their operation and be more cost-effective and safer, and so on, all the good words here.

But the take-away from this slide is

really the redlines here. We are trying to provide technical basis and guidance to address various issues. Sometimes what that means is, what's the technical basis for addressing a regulatory issue? That is what brings us here.

Now where we are right now, we have done a number of things over the years. Now you guys have the reports, some of the reports. There are some other scoping studies and sensitivity studies we have done in PRA that have not been published yet. Some of that will be published later this year. So we are continuing with that.

We are also working on additional guidance on protecting against common-cause failure. It has to do with failure analysis, and so on. That is also ongoing.

Also, we are working on better methods to estimate reliability of digital systems for the use of PRA. That is ongoing work.

Also, on better ways to do failure analysis for digital systems, there have been a number of cases where the utilities have come to us and said, "We put in this digital system and then it surprised us, and when we went back and looked, it turned out we didn't do a very good job in our initial failure

analysis before we put it in. We want help there."

So we are working on that. Stay tuned for that next year. We would love to come back and talk to you about these things when we are ready, but not yet.

We are planning to continue support of the NEI Working Group on Digital Issues, as appropriate, as we are asked, and so on.

Then there's this MOU, the Memorandum of Understanding that is now in place between EPRI and NRC Research. A number of areas are under discussion right now for continuing work, and Dan Santos can, of course, explain this more fully.

There's some examples here. More on operating experience. More on risk methods. More on diversity, well, actually, protecting against common-cause failure. Human factors has been suggested, and I guess there are some others. So that discussion is ongoing.

MR. HECHT: Ray?

MR. TOROK: Yes?

MR. HECHT: Could I ask a question with respect to the third bullet there, estimating digital system reliability based on design and process attributes?

MR. TOROK: Yes.

MR. HECHT: You are not going to be discussing that this time, but let me just ask you in general. Well, can you say anything about it?

MR. TOROK: Okay. Let me try to keep it at kind of a high level.

There are certain attributes in regard to digital system design in process that we believe makes the digital systems more robust and more reliable, more dependable, less likely to do bad stuff, all those kinds of things, right?

We can identify some of them. Like somebody said deterministic behaviors is a good one. You want to find that. These are the things we normally refer to as defensive measures, right? There are certain good design practices, and we believe that when you do a good job of implementing those, you improve reliability and dependability. So we are not being quantitative there. We are just saying we can kind of tell the difference between a good system and a bad system, right?

So what are the things we should be looking for, and how important are they? Now how do we take that into estimating reliability in terms of a number, if we have to do that?

So you end up talking about, where are your vulnerabilities and how good is your coverage relative to those vulnerabilities? That is the assessment you are trying to make.

It is not in an absolute world. It is in a reasonable assurance kind of world. So that is kind of what the discussion is.

The notion that ultimately you would like to be able to estimate reliability as a number is good enough for what you want to do in the PRA. Okay? And I'm not going to say we know how to estimate failure probability of a digital system. We don't know how to do that. But we know a lot about certain attributes that can make it better or worse, and that is really what we are going after.

MEMBER BLEY: Can I parrot something back to you and see if I'm catching what you are saying there?

It sounds like what you are doing is building a list of what one might call good practices for design of these systems, and then doing something akin to HazOp, or what's that other thing? PIRT. Trying to say, given this principle, if we implement that principle, how could we go wrong at a lower level while still meeting that principle at a high level?

Is that the kind of thing you're --

MR. TOROK: That kind of thing. It is getting a handle on, how good are we? What is our level of assurance here?

MEMBER BLEY: Just one last followup: are you publishing something on that soon? Or is this just --

MR. TOROK: Well, either late this year or next year.

MR. AUSTIN: Probably next year.

MR. TOROK: Probably next year, yes. Okay? And actually, there's going to be more discussion on that kind of thing later. Wait until Thuy gets up here and ask him more about that, okay?

No pressure, Thuy.

(Laughter.)

MR. HECHT: I am sorry. One of the things I didn't hear you say is I didn't hear you talk about the system architectures.

MR. TOROK: Oh, that's certainly a consideration, yes.

MR. HECHT: So, for example, one of the things that might be included in there is, are you going to be using an operating system kernel? What kind of device drivers are you using? What kind of

ports and port packages, and things like that?

MR. TOROK: That is certainly fair game for the discussion.

Did you want to remark now, Thuy?

MR. NGUYEN: Yes.

CHAIR APOSTOLAKIS: Microphone and identify yourself, please. Every single time you have to do that.

MR. NGUYEN: I am Thuy from EDF. I'm working in collaboration with EPRI.

To answer your question, this is a subject where we are very heavily involved. I'm from the Research Branch of EDF, and we do have research programs to try to determine, I would say, reasonable figures for failure engineering or beta factors for the digital systems, based on these deficiency measures and defensive measures.

For example, we do very deep analysis of the design of digital systems to determine whether the platform, the operating system, for example, could be or is less likely to be a cause of failure. If the operating system in the platform is, I would say, unlikely to be a cause of failure, then, for example, the beta factor could be lower.

With, I would say, the analysis of the

deficiency measures, we can determine what are, I would say, the main causes of failure in digital systems and focus our evaluation efforts on these , the main causes.

MR. TOROK: Now we don't have to publish the report, huh?

CHAIR APOSTOLAKIS: To what extent is EDF involved in your work?

MR. TOROK: EDF, Thuy is an EDF employee. In this case, they are the EPRI consultant as a principal investigator.

CHAIR APOSTOLAKIS: Are you using their operator experience, too?

MR. TOROK: That's a great question. Not yet. We're working on that. We're working on that.

CHAIR APOSTOLAKIS: Praise me, Ray. Praise me.

(Laughter.)

MR. TOROK: No, no, I hope you will encourage them to work with us here.

(Laughter.)

CHAIR APOSTOLAKIS: Mr. Nguyen is working with you, but the company itself is not collaborating --

MR. NGUYEN: No, no, we are developing --

we have the current ongoing project to analyze --

CHAIR APOSTOLAKIS: So that project is separate from what EPRI is doing?

MR. NGUYEN: Yes. Because although we publish in French, for example --

CHAIR APOSTOLAKIS: Those guys I know don't speak French.

(Laughter.)

I mean that's a great opportunity. You are such a large utility.

MR. NGUYEN: That's right.

CHAIR APOSTOLAKIS: Why don't you guys have a closer collaboration?

MR. AUSTIN: EDF is, of course, an EPRI member, one of our principal, larger, international members. In addition, EPRI and EDF have an MOU with EDF Research, where we do collaborate on items on a variety of subjects, materials, and we are looking for digital INC to make sure that we are leveraging each other's work as much as possible.

CHAIR APOSTOLAKIS: But the point is, can you come back in a year or so and say, "Now our operating experience includes the American experience and the French experience."?

MR. TOROK: We would love to be able to do

that.

CHAIR APOSTOLAKIS: Okay. So you're still wishing --

MR. NGUYEN: No, no, no. This will be done. This is currently being done. So we hope that we will have our, I would say, formalized analysis by the beginning of next year.

What we do have, I would say, informal results, but what we want to do now is to do it formally, going through all our database of significance events and to, I would say, have the same kind of --

CHAIR APOSTOLAKIS: You're operating similar to what EPRI is doing?

MR. NGUYEN: The approach will be similar, but the documents are different.

CHAIR APOSTOLAKIS: They are in French.

(Laughter.)

MR. NGUYEN: They are in French, and we don't have the same reporting mechanisms as in the U.S.

CHAIR APOSTOLAKIS: The same what?

MR. NGUYEN: Reporting.

CHAIR APOSTOLAKIS: Yes, yes.

MR. TOROK: But the good news is Thuy was

a participant in our evaluation of our OE. So he is very familiar with what we did. So that is going to help.

The other thing we should mention here is we are talking to another member with a lot of plants. This is South Korea, in regard to evaluating their data the same way. They have used digital systems for quite a while now. So we are hoping to expand the data we have.

CHAIR APOSTOLAKIS: Yes, very good.

MR. SIEBER: This is a key point that I tried to bring out earlier, that the experience doesn't really reside in the United States. In order to learn as much as we can, we need to engage ourselves with the French and --

CHAIR APOSTOLAKIS: The Koreans can help you very much in methods, too.

MR. AUSTIN: For Korea, that is a project we are starting now. We expect results late next year.

MEMBER STETKAR: This is probably a better question for Thuy, but I found that the processing and numerology of things that people call data often are not nearly as useful as the actual descriptions of what happened in a real power plant. That is why I

think your report is very, very useful. Those little snapshots in the appendix, which are backed up by more detailed descriptions are very, very useful.

I was curious whether the international experience is available at that level of detail or whether it is only going to be processed as we have done a study and where is our estimate of the failure rate.

MR. TOROK: No, no. It turns out, as I mentioned earlier, Thuy is very familiar with what we did. So he knows exactly what we are looking for in terms of how we evaluated the results.

As for the Koreans, we had the same discussion with them, and they are planning to send people to work with us, once they have gathered up some of their data, so that we basically make sure we treat their data the same way we did ours, to the extent we can.

MEMBER STETKAR: And when you say, "data", you mean actual descriptions --

MR. TOROK: Their descriptions and whatnot, that's right.

MEMBER STETKAR: -- reports of an event?

MR. TOROK: So they are going to identify -- they are going to go through all their events,

identify the key ones, and generate, basically, translations into English for particular events. Then we are going to work with them and our team, which means these guys you're looking at, to do evaluations of the same type we did here.

MEMBER STETKAR: I just wanted to make sure that you weren't entering into a collaborative nature, that you were just talking about, you know, "We did a study and here's my number," but you don't have access to actually the raw experience.

MR. TOROK: We want to be careful to treat their data the same way we treated ours. They are very interested in working with us to learn more about how we treated our data.

MEMBER STETKAR: The knowledge base is the important thing, what's happened in the real world. The processing of that is --

CHAIR APOSTOLAKIS: Can we move on now?

MR. TOROK: Yes. We are on a roll here.

(Laughter.)

Now I wanted to establish a context here that carries into the rest of the discussion. So this is an overview sort of sense here.

You know, what is EPRI doing? Why did we do this, that sort of thing?

Now when we started this work, it was really primarily in support of this NEI Working Group on digital issues. The digital issue that we were looking at was the one that some people call common-mode failure and some people call it common-cause failure, and some people call it defense in-depth and diversity.

For the purposes of our discussion today, all those things are the same. Okay? So we are going to talk about failure modes. We are going to talk about PRA, risk stuff, all more or less in that context of defense in-depth and diversity, things related to that. Okay?

And the reason we ended up working on those was because we were working with this NEI Working Group, and that was kind of the hot-button issue. So that is where the focus went at that time.

Now I would say there are lessons that go way beyond that particular context from this stuff, but that was where we started. So it is useful to look at what is out there in terms of guidance, NRC guidance, in regard to common-cause failure or defense in-depth and diversity.

This is a list. I tried to just list the documents that people always talk about. You always

hear these things thrown out in conversation.

The first one is the SECY-93-087 and the Staff Requirements Memorandum that goes with it. That goes back to 1993. Maybe it is getting a little old now. But, still, when people talk about policy, they are usually referring to that, the SECY and that SRM.

Now the next thing down is what is called Branch Technical Position 19. The full title is there, but it is often referred to as BTP-19. That I characterized -- this is just my characterization here -- that is what I call "what-to-do guidance" if you want to comply with the policy in the SECY.

Then, below that there's NUREG/CR-6303, which was a report from the early nineties generated by Lawrence Livermore. I characterize that as detailed guidance and technical basis. So, if you want to look for the technical basis, that is really where it resides, not in those other documents so much.

Then, after that, we talk about, people talk about the ISG, Interim Staff Guidance, and, in particular, ISG 2 is about defense in-depth and diversity. So people throw that one around.

Now ISG 2, it tended, I believe, to offer clarifications on the way the staff viewed those

documents above there on the list, the NUREG/CR, the BTP-19, and the SECY.

This is the one where the term, I think, high-occupancy vehicle lane comes in in terms of a way to expedite a regulatory review. It is also the one where the notion of the 30-minute criterion first comes into play. So that is the context of that one.

Now we know that, let's see, our guidance is based on the version of ISG 2 that was active, I guess, in 2007. Now we know that there is recently, in 2009, a modified version of ISG 2 that came out. I don't know if that is still considered a draft, but it has a different version of a 30-minute criterion.

Now our analysis is based on the earlier version. Later on, if you want, Dave can explain how the two different 30-minute criteria would play out in terms of his analysis results. So we can talk about that later. I just wanted to acknowledge that, yes, we know that we based our analysis on something that is now considered obsolete, I guess.

Then the last thing --

MEMBER BROWN: Wasn't that ISG 5 that Section 3 thing we talked about, as opposed to ISG --

MEMBER STETKAR: It is pervasive.

(Laughter.)

MEMBER BROWN: So they are leveraging it back into ISG 2? Because I hadn't seen anything on that.

MEMBER BLEY: My memory is ISG 2 had the 30-minute criteria. ISG 5 expanded on that.

MEMBER BROWN: Expanded it. Section 3 was what we --

MEMBER BLEY: Off of those alternative approaches.

MEMBER BROWN: Yes, and we had some comments on that that we fed back to the staff at that time. I haven't seen that that has been issued with the revised stuff in it. Am I correct on that? Okay.

MR. TOROK: The ISG 5 focuses on human factors and how long does it take for an operator to respond, and those kinds of things, as opposed to where the 30-minute criterion in ISG 2 is more about where might you need an automated diverse actuation system.

MEMBER BROWN: Yes, I understand that point. Okay?

MR. TOROK: Okay.

MEMBER BROWN: But I didn't know there was some variability on that definition that fed back into ISG 2.

MR. TOROK: Yes. And for our purposes, ISG 2 is the focal point, I suppose because we were doing the analysis before ISG 5 was written, right? So that is really why that one is listed here. You're right, though, ISG 5 does have some impact here.

Now the last thing, I characterize the staff positions, as somebody told me, staff position, that phrase has some legal meaning. Well, that may be true. If it is, I didn't mean that. I just meant substitute the word "opinions", if you like.

The document there is this SECY that came out earlier this year, SECY-09-0061, I guess. That is the one that is of interest to us because it includes comments on EPRI white papers.

Now there was an earlier letter in November of 2008 that, basically, had very similar comments. I think those, then, became the basis for the SECY.

So those are the things that are out there. Those are the context.

Now, if we move along here, we say, well, how does what we did relate to all those documents, and whatnot? As I said, we were looking at common-cause failure, defense in-depth, and diversity.

Our analysis really is centered toward one

part of the guidance that is out there. That means the guidance that refers to what you do in regard to postulated accidents and anticipated operational occurrences.

So there are two areas in regard to those things where our work is particularly relevant. The first one is in regard to the policy. Now I have oversimplified what the policy says there. I said the policy basically says identify your common-cause failure vulnerabilities and ensure that you've got adequate diversity for them.

Now it is pretty deterministic, pretty prescriptive. In fact, the SRM cautioned that SECY-93-087 was too prescriptive in some areas and shouldn't be taken too literally. So it is, basically, deterministic and prescriptive in regard to identifying vulnerabilities to common-cause failure, not in regard to assessing what adequate diversity is, because assessing adequate diversity is necessarily a qualitative engineering judgment kind of thing. It can't really be deterministic.

Now the other part of the guidance where our work really applies is in demonstrating compliance with the acceptance criteria of BTP-19, which tells you if you are okay relative to what the SECY is

looking for.

There are a number of -- what should I say? -- sub-items to that. I tried to characterize them here. In terms of acceptance criteria to BTP-19, you are looking at demonstrating adequate diversity or talking about corrective actions that are needed or providing a basis for taking no action.

Now I would say that, for the first part there, our OE and failure modes research really applies mostly when you are talking about identifying common-cause failure vulnerabilities. The risk insights apply mostly when you are looking at the acceptance criteria of BTP-19 because, even when we applied risk insights, we were deterministic about our CCF vulnerabilities. We weren't being risk-informed up there. We were being risk-informed in addressing the acceptance criteria.

Now, in that regard, our position was and our belief is that what we did in terms of approach and the results are consistent with current regulatory policy right now. So we don't necessarily see a need to change regulatory policy.

Now one thing I wanted to point out here, though, was that, if I go back to the document at the bottom, the SECY here, it mischaracterizes our

position and intent in kind of an interesting way. It offers an interpretation of what we were doing that we never intended. So perhaps we weren't clear about what we meant to do in two areas.

First, the SECY says that we were trying to use defensive measures to show that digital systems were not susceptible to CCF. That isn't what we were intending to do at all. All we were trying to do is use defensive measures and credit them in regard to assessing overall protection against common-cause failure, and this notion of trying to decide whether there's adequate protection against common-cause failure.

So I guess our position was, if you are serious about providing protection against common-cause failure, you really ought to be looking at these defensive measures because they are important. Any strategy for going after protection against common-cause failure ought to involve consideration of defensive measures and diversity where you need it, but they ought to be working together, and you ought to use whichever is better where it belongs. That was really where we were headed with this.

We weren't trying to show that defensive measures, if you had adequate defensive measures, that

the CCF, that you were not susceptible to CCF. That really wasn't the intent at all.

In fact, in some cases it was characterized as to show that CCF was not credible. Now I don't think you can necessarily do that with defensive measures, and you certainly can't do that with diversity either. But, together, you can do a pretty good job in terms of reasonable assurance. So that is where we were trying to go.

Now the other area that is --

MR. HECHT: Ray, can I ask a question?

MR. TOROK: Sure.

MR. HECHT: Defensive measures, as I inferred from the report, has at least three parts to it. One part of is what I would call the software development process quality.

MR. TOROK: Okay.

MR. HECHT: A second one is the -- how shall I say it? -- design features that one might include in the code. That would be, for example, things like don't use dynamic resource allocation; don't use --

MR. TOROK: Yes.

MR. HECHT: You know, do things to range-check your variables before using them, and things

like that.

MR. TOROK: Stuff like that.

MR. HECHT: Then, the third part would be some kind of overarching fault tolerance in the design.

Which do you mean?

MR. TOROK: All of those. And you mentioned architecture earlier. Architecture is a player here.

And you're on the right track here. What I would like to do is ask you to wait until we get to Thuy's talk and then bring that up again. Is that okay? Because that is certainly related to exactly what Thuy is going to talk about.

But the short answer is all of those things are part of it.

MR. HECHT: Because they have different implications, and they do different things.

MR. TOROK: Yes.

MR. HECHT: And I'm not sure that they should be grouped together.

MR. TOROK: I guess I would claim that they all are helpful in avoiding failures and common-cause failures. Well, some are maybe only one or the other. But, still, it is all part of establishing

assurance that you have adequate protection. They are all useful for that. Okay?

And some are much more useful than others. Let's wait until Thuy gets up here to try to pin him down more on that. Okay?

MR. HECHT: Okay. Is there any other connotation of what you called defensive measures that I didn't indicate?

MR. TOROK: I don't know. I thought he did a pretty good job of --

MR. NGUYEN: Yes. So you said the three legs, the first one is process, development process. I usually don't put that really in defensive measures.

MR. HECHT: Okay.

MR. NGUYEN: That is a given that we have to comply in every case.

The second leg is, I would say --

MR. HECHT: The design features in the code, that I might call robustness revisions.

MR. NGUYEN: Yes, fault avoidance. Fault avoidance is a very important approach for defensive measures.

MR. GEDDES: That could include hardware implementations as well.

MR. NGUYEN: And then three-elements.

MR. HECHT: Okay.

MR. TOROK: Okay? We have had a lot of discussions about this kind of thing. One, oh, I hate to bring this up in a way, but if we were talking about cars, everybody knows cars, right? We would say, well, cars don't stay right-side-up most of the time because they have a good software development process or a good design process. They stay right-side-up because they have four wheels that are pretty far apart and a low center of gravity. Those are design features that add protection, right?

It is the same kind of thing you are talking about here in the digital system. So there are process attributes that are good, but there are also design attributes that are very important. You don't want to forget about those. That's all.

So we will go back to that theme, actually, over and over again. So you will have plenty more chances to comment on that.

MR. HECHT: I am sorry.

CHAIR APOSTOLAKIS: Don't be sorry. Don't be sorry.

MR. HECHT: Okay. I heard Thuy say that process wasn't part of it, and I heard you say that it was.

MR. TOROK: Well, yes, and that's sort of an ongoing discussion.

CHAIR APOSTOLAKIS: That may be --

MR. TOROK: Frankly, I am not sure it matters.

CHAIR APOSTOLAKIS: It doesn't matter.

MR. TOROK: They are both good things.

CHAIR APOSTOLAKIS: But, Ray, you have been into it for 55 minutes, and I still haven't seen a single operating experience.

MR. TOROK: Well, I feel the same way.

(Laughter.)

I'll tell you what. I think you guys --

CHAIR APOSTOLAKIS: Why don't we speed up the thing?

MR. TOROK: Great. And you guys scheduled your first break, I think, for 9:30, right?

CHAIR APOSTOLAKIS: No, it will be when I say it will be.

(Laughter.)

MR. TOROK: Okay, fine. So let's try to get to the good stuff, okay?

MEMBER BROWN: But before you do that, I can't stand this anymore.

(Laughter.)

We brought up the fault tolerance issue. While I agree fault tolerance is very, very important, if you step it up at a high level, you see there's fundamentally two types of systems. We've got feedback control systems. We control turbines. We control the feedwater system, the blah, blah, blah.

We also have what I call once-through systems. You measure things. You decide I'm going to shut it down and stop everything right now. The feedback is put the rods on the bottom, whatever the control devices are, or jack them up, whichever direction they are going to go.

So, when you talk fault tolerance, you really have to look at the application of the systems and decide what type of fault tolerance you are looking for.

MR. TOROK: Yes.

MEMBER BROWN: If you look at the feedback control systems, there's one type. You would like to keep systems on the line. You don't want the failure of one thing to all of a sudden dump stuff offline. So there is a basis for some approaches to the fault tolerance that you take in those systems that are not very useful in what I call the once-through systems, where you want to shut it down.

For instance, data exchanges between theoretical, independent channels. In a feedback control system, data exchange can help you, if it is designed properly. In a once-through system, data exchange leads to reduction or compromise of independence, where a single thing happening in one, where you have exchanged the data, and now it goes over and it compromises all four other channels or three other channels, whatever the numbers are. All of a sudden, you don't have a protection function.

So, when somebody starts talking fault tolerance and data exchanges, antenna go up. There's different ways of data exchange. Whether it is sensor data or whether it is output trip data, or whatever it is, data exchanges between protection-type or safeguards-type channels can be very detrimental to your ability to say I meet my requirements.

MR. TOROK: I think we agree with you on that. In fact, I think the way Thuy might say that is there's no magic list of defensive measures that applies everywhere. It depends on the context.

MEMBER BROWN: I just wanted to frame the fault tolerance into sort of out of what I call the more academic, you know, beta factors, and all this other kind of stuff, because I don't understand that

kind of stuff. I just look at stuff either works or it doesn't.

MR. TOROK: And you're right, often when you start talking about these things on a theoretical basis, it is good to pull yourself back into real life once in a while and think about that.

MR. SIEBER: I think I don't want to take up a lot of time, but there was something that was said that I think is vitally important. In INC, as in other branches of engineering, people tend to put themselves in boxes.

One would say that, regardless of the system and its components' behavior, we could build an INC system that will operate it. On the other hand, we would not have a lot of digital INC applications to three-element feedwater control, for example, if they would design the float control valve properly.

(Laughter.)

I think, in order to have simple, reliable systems, you have to pay attention to the dynamic characteristics of the devices that it is controlling. I don't see that in any of this, other than a tacit recognition that maybe that is the case.

I think you really have to look at things like valve operators and other actuating devices in

order to make the system work right. It doesn't make any difference whether it is analog or digital; if the operating components don't work right, it's not going to be a success.

MR. GEDDES: And we've seen OE where a digital feedwater system, for example, would go in, and it is the final control element that really had the problem.

MR. SIEBER: Yes.

MR. GEDDES: And the digital system didn't anticipate a poor implementation of the final control element.

MR. SIEBER: Or it could actually make it worse.

MR. GEDDES: Or reveal the condition.

MR. TOROK: That's right. Exactly that.

CHAIR APOSTOLAKIS: Okay, let's move on.

MR. TOROK: Okay. Let's try to move on, right?

CHAIR APOSTOLAKIS: Yes.

MR. TOROK: Okay. There's only one other point I wanted to make here.

CHAIR APOSTOLAKIS: Now can you go to slide seven, Ray? Start talking about operating experience.

MR. TOROK: Yes, sir.

CHAIR APOSTOLAKIS: I am sure you will find opportunities to interject your thoughts.

MR. TOROK: Okay. Okay.

CHAIR APOSTOLAKIS: Seven. That's it.

MR. TOROK: Okay. This is the overview of the key points in the various areas. Okay? All this stuff on one slide, in the world of operating experience, the results of our studies is basically software has not been any more problematic than other contributors to common-cause failure. We have seen evidence of actual, well, of potential and actual software common-cause failures, but a lot more of other kinds.

The recommendation was software is doing pretty well. What we need to do is figure out systematically why it is doing that, and make sure we capture that knowledge and continue to apply it.

Now, in terms of failure modes, it is not as simple as talking about just failure modes. There are failure mechanisms, which produce failure modes, which have effects on systems. In analyzing events, as someone using PRA, the important thing is to understand what you care about relative to mechanisms, modes, and effects.

We are going to talk more about that later, but the main point is the PRAs don't necessarily need low-level failure mechanism knowledge to model what they model and to generate risk insight.

They do need to have other things nailed, though, in terms of effects and some failure modes, and so on. We will have a lot more for that later.

Let's see, now in terms of protecting against common-cause failure, both prevention and mitigation are really important. Okay? And it is not just one or the other. There again, we will have more about that later.

As far as the PRA insights go, we believe insights are possible today, generating real insights today, using existing techniques. I am going to show you what we did that makes us think that.

We are thinking in the use of PRA, where it is appropriate, and PRA does a good job of figuring out, I think, whether its results are appropriate, and so on. But where it is, we should be doing more of it. We should use it. We should take advantage of it.

Where a PRA is really nice is if you are talking about we said like contributing adequate diversity, for example. It is a subjective thing.

PRA risk insights can be helpful in making that judgment. That's all.

Now what we are looking for in terms of coming to this group, we would appreciate your concurrence in that we are aimed in the right direction here. The first thing there, I don't think anybody argues with. Continue to gather and apply OE lessons on failure causes, corrective actions, and preventive measures.

But maybe the more important thing here is this notion that, when you are evaluating these events, which I say you have to be really careful about how you define things and how you break things down, and get some common understanding of what is a reasonable way to do that.

We have taken a shot at it, and we will explain to you how we shot at it. There are other ways you could do it. So that is, I think, an ongoing issue.

In terms of crediting defensive measures, we think defensive measures are really important in terms of protecting against common-cause failure, and those should be pushed.

As far as risk methods go, what I just said --

CHAIR APOSTOLAKIS: Ray, let me interrupt you.

MR. TOROK: Oh-oh.

CHAIR APOSTOLAKIS: I don't know how many times you have come before the ACRS, but you're giving us motherhood statements here.

MR. TOROK: Okay.

CHAIR APOSTOLAKIS: Can't you move on to the real thing?

MR. TOROK: I'm sorry.

CHAIR APOSTOLAKIS: I'm sorry to interrupt you, but we are behind. I mean we know what we should be doing.

MR. TOROK: Okay.

CHAIR APOSTOLAKIS: The question is how to do it.

MR. TOROK: Okay. I'm so sorry.

CHAIR APOSTOLAKIS: I hope you don't take this the wrong way, but we really have to get moving here.

MR. TOROK: No, you're right.

CHAIR APOSTOLAKIS: Okay.

MR. TOROK: You're right. Okay.

So some review here is we looked at 322 events. Our focus was on --

CHAIR APOSTOLAKIS: That's good. Now you're talking.

MR. TOROK: -- common-cause failure.

CHAIR APOSTOLAKIS: Now you're talking.

(Laughter.)

MR. TOROK: Okay. Well, I'm going to try here, okay?

CHAIR APOSTOLAKIS: Great.

MR. TOROK: Now we said, actually, potential common-cause failures, you care about both of them. We are not saying one is important and the other is not. We care about both of them. So we are trying to find that.

It is useful to note that we are only looking at bad stuff here. The success stories were not addressed. So you are going to see some reports of common-cause failures in feedwater systems, but you are not going to see any reports of successes there.

CHAIR APOSTOLAKIS: As I was reading your report and this statement, "Look for actual and potential CCFs," it triggered my memory. As you know, or perhaps you know, there was a major joint effort between EPRI and the NRC back in the eighties to look at hardware common-cause failures. They faced the same problem, actual and potential.

They came up with some diagrams that helped them process the potential common-cause failures and produced some results that could be used in PRA applications.

I am not saying you should be doing that, but it seems to me it would be helpful if you went back and said, how did these guys do it? Maybe we can do something similar or modify it to our problems.

MR. TOROK: We should look at that.

CHAIR APOSTOLAKIS: Yes. They have some nice little diagrams. You know, if I had three trains but only two had been affected by the CCF, if I looked into the details of what happened, what is the probability that the third train could have been involved? You know, those kinds of --

MR. TOROK: Yes. Okay. Yes, we will look at that. We will look at that.

MEMBER STETKAR: I am going to jump ahead here a little bit because, unfortunately, I need to get out. But it actually dovetails with something George just said.

CHAIR APOSTOLAKIS: Go ahead.

MEMBER STETKAR: Do you have a copy of the report available that you can bring up on the screen?

MR. TOROK: Yes.

MEMBER STETKAR: Okay. When I was reading through the report, and I do think it is important to have a clear understanding of the definitions and the classifications --

MR. TOROK: Wrong one.

MEMBER STETKAR: I don't think it is ACRS's function to go back and review all 322 screen captures in the appendix. But one event, in particular, was called out. It is called out as a good example, would be an event -- this is an example of a software failure. It is event No. 17, in particular. I wanted to pull up the screen.

MR. AUSTIN: A screenshot of the event?

MEMBER STETKAR: A screenshot of the event.

MR. AUSTIN: Okay.

MEMBER STETKAR: Only because I want to use that to try to understand how you did your classification process because it is important for us to understand how you thought about that classification when you start presenting the results of all of your classifications.

MR. TOROK: So did you want to see the --

MEMBER STETKAR: I would like to see the screenshot of event No. 17, if you can pull that up.

MR. TOROK: Okay. From the appendix of the report?

MEMBER STETKAR: Yes, sir.

It's on page No. 77 in the PDF file.

It's page 19 of the appendix.

MR. TOROK: What was the number?

MEMBER STETKAR: Nineteen of the appendix or 77 of the PDF file, depending on what you're --

CHAIR APOSTOLAKIS: Are you there?

MR. TOROK: Yes, just about. I'm sorry, what was the event number?

MEMBER STETKAR: Seventeen. There it is.

MR. TOROK: That is event 17.

MEMBER STETKAR: Yes. You probably can't read it there, but this event apparently happened, as best as I can tell -- it is characterized in the text on -- I have to jump back and forth here in my own file -- as an example of -- bear with me here --

CHAIR APOSTOLAKIS: Can you give us a short description of the event? Not everybody is up with it.

MEMBER STETKAR: Yes. The event was, apparently, a change that was made to the software for a digital feedwater control system. Because of inadequate verification and validation of the

software, there was a logic failure that caused a main feedwater pump speed to go to zero, and it caused a reactor trip because of the loss of feedwater, and the rapid trip occurred at 30 percent power, which you can probably read up there.

MR. TOROK: Right.

MEMBER STETKAR: Now the event summary, if you look at the checkboxes, and if you read the little blurb under there, it says, "Affected master controller in one train, not a CCF." So the little checkboxes for CCF and potential CCF are not checked off.

Now my question is I need to understand the thought process for doing this classification because, if this is a software verification and validation failure in terms of implementing a new set of software that happened to cause a trip of a single feedwater pump, and, of course, at 30 percent power, you only have one feedwater pump running. So it could have only affected that one pump.

Why is that event not at least a potential common-cause failure? Or in George's construct, why isn't there any probability that it might have been a potential common-cause failure.

I'm bringing up -- and I wouldn't have

looked at this event except that it was highlighted in your text as a good example of a software failure.

CHAIR APOSTOLAKIS: Is there a discussion though? I don't remember.

MEMBER STETKAR: No.

CHAIR APOSTOLAKIS: Arguing why it is not?

MEMBER STETKAR: No, no, no.

CHAIR APOSTOLAKIS: Okay.

MEMBER STETKAR: No. It's only highlighted in the report as an example of what a software failure is in the context of software failure.

MEMBER BROWN: Well, that was back in the text, not in the appendix, wasn't it?

MEMBER STETKAR: But when I read the event, I said, oh, okay, this is a software failure. I understand it is a software failure, but why aren't the checkboxes checked off and why, in particular, is it specifically stated that it is not a common-cause failure?

I think that is important. It goes back to what I was saying a little bit earlier. Individual analysts' interpretation of these events and classification of the event may be subject to discussion. However, the event itself, the

description of the event itself, should not be --

MEMBER BLEY: Can I turn your question around a little?

MEMBER STETKAR: Sure.

MEMBER BLEY: Would you classify this the way it is marked? Was this an error or is this the way you would have classified it?

CHAIR APOSTOLAKIS: This may be an oversight. Would you still do it this way?

MEMBER STETKAR: I am assuming there is more information behind this. This is only a single screenshot.

MEMBER BLEY: And maybe you can't answer it on the spot here.

MEMBER BROWN: Wasn't that one discussed in the text of the thing as far as discussing it?

MEMBER STETKAR: The only reason I looked at it, it was discussed in the text as -- they were discussing the different classifications, and they said, well, some events are classified as software failures, and, for example, go look at event No. 17 as one of those events. That is the only reason I went to look at it. I certainly didn't review 122 screenshots. I looked at a number of them. I didn't look at all of them. I swear to God I didn't look at

all of them.

(Laughter.)

CHAIR APOSTOLAKIS: I believe it was event 220 that was the only common-cause failure or I think --

MEMBER BROWN: But 222 --

CHAIR APOSTOLAKIS: Two twenty-two.

MEMBER BROWN: In the 1E systems. Is this a 1E? I presume this is a 1E system?

MR. TOROK: No.

MEMBER BROWN: Then we've got to go back to the next section.

MR. TOROK: Is it event 17?

MEMBER STETKAR: It is event 17. If you look back at the non-1E summaries, it is not listed in the non-1E -- you know, there's a good cross-reference. I really like the way the report is put together.

MR. GEDDES: Mr. Stetkar, can you tell us where in the report the reference is made, so we can get to it quickly?

MEMBER STETKAR: Yes, it is on page -- hold on a second because I just made notes based on PDF file page --

MR. TOROK: That's okay. We have that.

MEMBER STETKAR: It's on page 3-2 of the report.

MR. TOROK: Oh, okay.

MEMBER STETKAR: The first paragraph on page 3-2. In the PDF file, it is page No. 28. It is page 3-2 of the actual report, and it is the first paragraph.

And the quote, I mean I can read it because nobody can see it. It says --

MR. TOROK: Okay, it's up here now, I think.

MEMBER STETKAR: "Events involving digital technology mishaps are referred to in the report as, `software events' or `software failures'. A good example would be an event caused by a fault or bug in a software control algorithm which was then missed during V&V and testing (example: event 17 in Appendix A). This would be considered a software design error and an indicator of potential weaknesses in the process used for software development."

MR. TOROK: Right.

MR. AUSTIN: So the question is, when you went back and looked at that, it is a software error, but why was it not classified as a CCF?

MEMBER STETKAR: That is either a real CCF

or a potential CCF.

MR. GEDDES: Yes. I think in that case, that checkbox marked potential CCF either at the subsystem or the system level should have been checked.

MEMBER STETKAR: But, see, whoever is doing the analysis actually wrote in the little explanatory box, "not a CCF".

MR. GEDDES: Right.

MEMBER STETKAR: I mean somebody --

MR. GEDDES: That was me.

(Laughter.)

MEMBER STETKAR: Okay.

MR. GEDDES: We included these text boxes in the final report because it is a form of commentary. Okay? It is just commentary.

MEMBER STETKAR: Yes.

MR. GEDDES: The taxonomy and the classification scheme, we have detailed slides. Dr. Apostolakis, the idea of using figures to diagram what these terms mean occurred to us after our last appearance here, and we can show you exactly what we mean by these checkboxes. Okay?

Now, in this case, event 17, should have checked off potential CCF either at the system or

subsystem level. If it is a single train, in other words, one of the feed pumps might have had this defect, and another one didn't. I have to go back and read the event. It's been two years since we looked at all the event details.

MEMBER STETKAR: No, that's why. It's a bit unfair just to take it out of context because of the screenshot.

MR. GEDDES: No, it's fair because we are here to represent this information. Okay?

CHAIR APOSTOLAKIS: But they enjoy it.

(Laughter.)

MEMBER BROWN: No. I was going to -- let Bruce finish.

CHAIR APOSTOLAKIS: The issue, as you know, you seem to be familiar with it, is always, what did we learn from what we see? I mean, if it is a straight common-cause failure, okay, I mean it is an unfortunate occurrence, but from the analysts' perspective, it is not very challenging.

But when you have these situations where something happens in one train or two trains, and then you have this other train over there, and you have to go deeper into what happened in order to make some inference as to its potential applicability to the

other train, I think that is the challenging part.

I think those little diagrams that I was not very impressed by at the time seem to help because, first of all, they tried to establish some consistency among analysts because everybody is using the same diagram.

But, also, there will be a little box there, you know, what is the condition of probability that they could have propagated? So then you are forced to think about it, which also forces you to think back about the applicability to the other train.

So I think that would be a very useful thing to revisit. It is a whole series of reports, as I remember. PRG was involved at that time, and it was joint EPRI/NRC.

MEMBER BLEY: And Idaho was in that.

CHAIR APOSTOLAKIS: I'm sorry?

MEMBER BLEY: Idaho was in that, I think, National Lab.

CHAIR APOSTOLAKIS: Idaho was involved, but Ali Mosleh was --

MEMBER BLEY: Bruce, you were about to say something about the text field.

MR. GEDDES: Yes.

MEMBER BLEY: I wanted to hear what you

were about to say.

MR. GEDDES: Well, the text field is a commentary. We took notes as we went along. We compared and contrasted. We went through hours and hours and hours of collaborative review of some of these event reports.

This one was interesting from how the error occurred and how it propagated into the system. That is why we called it out in the text. But we focused probably most of our energy on the common defects, especially on 1E systems, the 49 events.

We ended up converging on a certain meaning of these checkboxes that are in the top of the figure, and scrubbing that taxonomy and usage very, very carefully. Okay?

I think, for this event, that checkbox should have been checked. I would have to go back and read the event report and try to reconstruct how we got here.

MEMBER BLEY: Okay.

MR. GEDDES: But it does describe the condition that would -- without reading the report again, I would suppose that in this case there's probably redundant controllers on each feed pump, and the logic would be incorrect in both controllers.

This refers to a master controller. So that architecture, I would have to go back and read it.

MEMBER STETKAR: You have to look at the individual --

MR. GEDDES: You have to read it, yes.

MEMBER STETKAR: Because it is a master controller, and you don't know how it works.

MR. GEDDES: Without reading the report itself, and I've got a total screenshot of the whole database. We redacted certain portions because it reveals OE numbers from INPO sources that they are not real comfortable with making public. Okay? I've got the plant name and all the details here.

MEMBER STETKAR: Yes, yes, yes.

MR. GEDDES: I can go back and read the report. There might be a good reason why I didn't check the box.

MEMBER STETKAR: There might be.

MR. GEDDES: But I don't recall.

MEMBER STETKAR: There might very well be. As I said, it is unfair in this forum to put you on the spot just because of that restrictive screenshot.

MEMBER BLEY: On the other hand, this is the guts of what everybody is worried about.

MEMBER STETKAR: Right.

MEMBER BLEY: So anything that is a little questionable, I would hope, when you are doing the analysis, you would make a big effort to explain just why the heck this isn't what it looks like it is.

MR. GEDDES: Right.

MEMBER BLEY: So it's fair or not fair, but if we miss just a couple of these --

MEMBER STETKAR: Yes, there aren't many.

MEMBER BLEY: -- we miss the whole story, and we get the wrong impression here.

MEMBER BROWN: There was another one that, if you look at page 4-7, there was one where they are talking about one potential one. It springboards from what John brought up.

"Five of eight automatic self-test routines running in each of four asynchronous sequencer channels had an error in the application logic that would have prevented an actual safety injection signal from passing through while in auto-test mode."

Well, it is kind of an interesting thing because you have to balance all these. I mean this whole issue of independence, reliance on self-testing, data interchanges, as soon as you start doing data

interchanges, if you've got self-testing routines that do that, I mean you can screw everything up. I mean there's all these benefits, but you can nail yourself to the wall.

So this idea of whether it is a CCF or whether it is design issue, you know --

MR. GEDDES: Can I make a suggestion? We have, first, slides that pictorially or graphically describe what these terms mean. That would be probably helpful --

MEMBER STETKAR: That would be great.

MR. GEDDES: -- to get that first. Then we can look at certain events.

We brought several back-up slides or we even have that particular event in the main body of our presentation. We spoke of it last time, and Dave Blanchard did some significance determination evaluation on it. I would like to invite him up when we get to that event.

So, if we can get the taxonomy clear first, which is part of the commentary, I think, then we can look at the events all from the same frame of reference.

MEMBER BLEY: I have just one other question. When George talked about those diagrams,

and you said, when you guys left here, you were thinking about them --

MR. GEDDES: Yes.

MEMBER BLEY: -- have you actually done anything with that idea?

MR. GEDDES: Well, yes, that is what we are prepared to show you. Maybe not the same kind of diagram, but we drew pictures to describe what we mean.

MEMBER BLEY: Okay. So you are going to show us that?

MR. GEDDES: Yes. Yes.

MEMBER BLEY: Okay. I was looking at a table, and that's great.

CHAIR APOSTOLAKIS: Would this be a good time to take a break?

MR. TOROK: Sure.

MR. AUSTIN: If you think so.

(Laughter.)

CHAIR APOSTOLAKIS: Fifteen minutes.

MEMBER BROWN: Sorry. What did you say?

CHAIR APOSTOLAKIS: Fifteen.

MEMBER BROWN: Oh, okay, I thought you said 10.

(Whereupon, the foregoing matter went off

the record at 9:50 a.m. and resumed at 10:09 a.m.)

CHAIR APOSTOLAKIS: Okay, we are back in session.

MR. TOROK: Very good. So we are picking up here with regard to our first topic, operating experience review. You guys all know we looked at 322 events and all that.

The main point here is I just wanted to point out that, when we were before you back in March and April last year, we had a white paper on the subject. We since published this final report, which you all have. That report was published in December last year. We sent it into ACRS and NRC in January.

It expanded the discussion of various things, in part, to address comments that you guys raised, as a matter of fact. Then we added this appendix in the back that had the brief descriptions of all the events that we have been talking about.

I wanted to mention that, in evaluating these events, one guy, Bruce didn't just decide everything. We had very detailed discussions involving several people, in fact, anybody who cared to comment almost, but some really, I want to say, heated discussions in regard to what the event descriptions really meant among some various experts.

Bruce, obviously; Dave Blanchard, Thuy participated, Vick Fregonese, sitting here, from REBA, he participated, and there were some others. So it was a panel. We had some very interesting discussions.

We would agree that some of this stuff is subjective, and a different group might arrive at somewhat different conclusions. You are seeing some of that stuff today. Okay?

MR. GEDDES: It would have been good, I think, to interact with staff on some of these.

MR. TOROK: Well, actually, we --

CHAIR APOSTOLAKIS: Which you will. You will at some point.

MR. TOROK: Well, we actually invited staff on a number of occasions, but their restrictions prevented them from discussing it with us.

CHAIR APOSTOLAKIS: Under the MOU, that will not be the case?

MR. SANTOS: Dan Santos.

The answer is, yes, we do plan to collaborate in the future.

CHAIR APOSTOLAKIS: Good.

MR. TOROK: Yes, it would have been great to have them more involved in these discussions.

Anyway, moving right along, key terms, I am not going to talk about all these. These are here for reference. It is a list out of the report. Bruce is going to actually explain what some of these things mean in a pictorial form in a minute.

Before we go on, there are two of them that I did want to talk about a little bit because every time we do this presentation the same thing comes up.

One of them is, what's an event? I think Myron raised this question the last time here. For us, the purpose of our evaluation, a digital event was basically anything that involved or affected a digital system and was reported. Okay? It's not necessarily a plant transient or an accident or anything like that. It had to do with what was available in the reports.

MR. GEDDES: I would like to point out we used keyword searches like software, digital, computer, feedwater, protection, control. We tried to cast as wide a net as possible, and we didn't exclude any data.

So we brought out all 322 events and then systematically reduced it to the most interesting ones.

CHAIR APOSTOLAKIS: These data came from --

MR. TOROK: LER database.

MR. GEDDES: Well, we used the INPO OE search engine, which included EPIX and other sources.

MR. TOROK: And then the NRC database with the LERs. So those are the main two places.

Actually, Mike Waterman shared with us a list of events that he had been collecting over a number of years. I think there were 340-some-odd ones of those. It was just a very brief list of events.

We looked for the writeups on those events, so that we could include them. Of those 340-some, we actually found 160, and we included those. On the others, since we didn't have a detailed description, there wasn't anything we could do with them.

MR. GEDDES: We called Mike and said, you know, there's a slew of events on this list that we couldn't find the source documents for. We did a search, but couldn't quite get all of them.

MR. TOROK: Right. Okay. So that's what events meant to us. So don't read any more into it than that.

Now the other key term that keeps coming

back at us, and for a good reason, I think, is this notion of a software event. Frankly, the definition that we have here, which is the one in the report, does not do justice to what we really did. So I wanted to explain that a little more.

It is just events involving design defects. It is really somewhat broader than that. What we were trying to get at was this notion that it was events that involved digital behaviors of the system one way or another. So it was broader than just design of the software.

But where digital aspects didn't really play a role, we didn't include that. For example, a setpoint error, well, the setpoint errors can be done in analog or digital. They don't care which the system is. So that wouldn't be a digital-specific event for us or what we would call a software event.

However, if there were a bug in the code, let's say, in the end that was missed in V&V and testing, and so on, that would be a software event, regardless of where it actually came from. Did it come from a programming error? Did it come from something they missed in the testing? Sometimes you can't tell. Sometimes you can. Right?

Now there were other cases where, if the

event resulted from a misunderstanding of the basic phenomenon, such that the digital system did not recognize an event it was supposed to have recognized because the requirements were wrong way back at the beginning, that probably isn't something we would call a digital event because an analog system based on those requirements would have the same problem.

MR. GEDDES: Software error.

MR. TOROK: A software event. I'm sorry. Yes, we wouldn't call it a software event.

We actually have one example of that where the problem was caused by the fact that the requirements did not anticipate an actual behavior of the plant, such that the system didn't recognize it when it happened.

MEMBER BLEY: I don't remember, did you keep those in a separate class? Because that would be an interesting class to look at.

MR. GEDDES: Yes. In our Pareto charts, we separate all those out and we show how many there were and what the distribution looked like.

MR. TOROK: Let me show you that momentarily. Okay?

That's all I wanted to say about those things. Are there any questions about that? Are we

okay on that, on those definitions right now? Those are really key.

Another example was, if it was a parameter error, we probably wouldn't call that a digital -- or a software event.

Now if it was something where the software design introduced tremendous complexity because it was doing things that a comparable analog system was unable to do, we would call that one a software event. Okay?

But, obviously, there is some judgment involved here. We argued among ourselves about how to do it.

Anyway, so I just wanted to lay that groundwork.

Now, finally, ladies and gentlemen, the rest of those definitions on that list --

MEMBER BROWN: I'm sorry, I do have a question. I didn't ask it.

How do we cause failure relative to software? Do you all view that as a piece of software that gets corrupted and then can propagate to cause a failure in other channels, or whatever, based on its propagation? There's design errors where you make a design software error. Now that can be a common-cause

failure, but it is not a software corruption issue. It is literally the programming error in terms of how you execute.

MR. GEDDES: A latent failure.

MEMBER BROWN: Yes, exactly. Yes, you can't always find those by testing or any other thing.

Do you all differentiate, is the question, or do you even recognize what I said?

MR. GEDDES: I think I understand what you said. Of course, there's data corruption, and then binary, like memory errors, can affect the way software is supposed to behave in a system where the software is loaded and running. Okay?

But, in this context, we talk about design defects. That is what we mean.

MEMBER BROWN: Okay. All right.

MR. GEDDES: In other words, the software is properly loaded, you know, and it is a latent defect introduced in the design process.

MEMBER BLEY: The other one that Charlie is talking about, sometimes register overflows, something happens and corrupts the code, how do we find those in your data?

MR. GEDDES: Well, there's only a handful. There's very few. Okay?

MEMBER BLEY: But, conceptually, there have been examples in other areas where those have been disastrous.

MR. GEDDES: Right. I tagged those in the data. I can show you where they are.

MEMBER BLEY: Okay. Do they have a name? Or you will show us when you get there?

MR. GEDDES: I'll show you.

MEMBER BLEY: Okay.

MR. TOROK: Okay?

MEMBER BROWN: Okay. Yes.

MR. TOROK: Now this picture is where we address the rest of the definitions, and Bruce is going to take it here, please.

(Laughter.)

MR. GEDDES: Good morning.

(Laughter.)

We prepared this chart after one of our last appearances with the ACRS, and we appreciate the feedback.

This is a two-channel construct. It could be a four-channel construct. We don't mean to exclude those.

But, in the sense of redundant and independent, potentially independent, control systems

or non-safety systems lose their independence at one point or another. Especially where there's one final control element, a sense of independence gets lost. That is where we get into master-slave architecture, that sort of thing.

But let's take this from a simple point of view. In a two-channel system, across the middle of the chart, we have the notion of a common defect, concurrent triggers, and whether or not there was a failure.

So, for a common-cause failure, there's two ingredients, a common defect and a concurrent trigger. Okay? We made that distinction, we came to that distinction while we were analyzing the data, especially the hours we spent going over and over and over the events that reported a common defect.

The software, by definition, where it resides on multiple channels or trains, is a common defect, if it has a defect.

MEMBER BROWN: A latent defect?

MR. GEDDES: Correct.

MEMBER BROWN: Okay.

MR. GEDDES: A design defect or a latent defect.

So the first column, where we see there's

no common defect, no concurrent triggers, no failure, both channels are green. That means no defects, no problem.

We have had some events where the digital system was mentioned, but it wasn't really part of the problem at all. It just happened to be nearby. Maybe it was a valve problem, and the digital system responded appropriately.

The next column would be a single failure in which there would be no common defect, but a single failure. It could be, typically, a hardware problem. Channel one of this construct would be failed.

The next column over, where we see now we have the presence of a common defect, and we use the yellow box to show the presence of a defect, and the dotted red line to show where the triggering condition might be.

We saw some events where there was software in multiple channels. The defect would be triggered by a sensor failure, for example. In a deterministic world -- and I'm an I&C design person; I'm not a PRA guy -- I view that deterministically as a single random failure that would not propagate into a common-cause failure. Okay?

It would have to take concurrent random

failures of multiple sensors in that example to result in a potential or an actual CCF.

CHAIR APOSTOLAKIS: Are you under the --

MEMBER BROWN: There's an external cause then?

MR. GEDDES: External trigger.

MEMBER BROWN: External trigger. I'm sorry. Thank you. External trigger.

I am sorry, George.

CHAIR APOSTOLAKIS: Are you under the third column?

MR. GEDDES: Yes.

CHAIR APOSTOLAKIS: So you're saying on channel one there was a defect and a trigger?

MR. GEDDES: Correct.

CHAIR APOSTOLAKIS: And channel two had a defect?

MR. GEDDES: The same defect, but --

CHAIR APOSTOLAKIS: And it is common, right?

MEMBER BROWN: The trigger is unique.

MR. GEDDES: It's one defect, and the --

CHAIR APOSTOLAKIS: It was a potential common-cause failure, I suppose?

MR. GEDDES: We make the distinction about

the nature of the trigger that can create the failure. If the trigger, for example, is due to a sensor failure, which would be a random failure, we draw that distinction and say, deterministically, only one channel can fail at a time due to a sensor failure.

CHAIR APOSTOLAKIS: Even if the defect is common?

MR. GEDDES: Yes. Yes, that's a very, very important distinction that we make.

MEMBER BLEY: I see the logic of this, and I see it is useful. There's something that is a little unsettling and doesn't quite go to the thing George was talking about earlier.

Once you get two yellow boxes here, you've got a common defect. This event might not have had a trigger, and there might not have been a common trigger, and there might not have been a failure. But right when you see the common defect, it seems to me a place where you ought to start thinking, are there any triggers out there that could have led to the kind of problems we are worried about?

MR. GEDDES: Yes, and that is what we looked for. That is what we evaluated.

CHAIR APOSTOLAKIS: It is a latent common-cause failure of the system.

MEMBER BLEY: It is, and labeled no common cause --

MEMBER BROWN: It is a design error in the system, but it has to be triggered from an external source.

CHAIR APOSTOLAKIS: I don't understand what --

MEMBER BROWN: If it hasn't been --

MEMBER BLEY: But we have found it.

CHAIR APOSTOLAKIS: But why isn't that a potential common-cause failure?

MEMBER BROWN: Because you only have one trigger, it is a single failure that doesn't trigger both --

CHAIR APOSTOLAKIS: This is one that actually happened.

MEMBER BLEY: Yes.

CHAIR APOSTOLAKIS: In another situation, you might have a trigger that --

MR. TOROK: It comes back to this.

CHAIR APOSTOLAKIS: Yes, go ahead.

MR. GEDDES: Because of this very discussion, which we went round and round and round amongst our peers, "Concurrent trigger. Triggers which occurs over a time interval sufficiently short

that it is not plausible that resulting failures due to a common defect would be corrected."

In other words, a single sensor failure would reveal itself before, and we could correct it before another sensor could fail in that context. That's what we mean by that, that we could discover and correct a condition before it would propagate into multiple channels.

CHAIR APOSTOLAKIS: But are you sure that there are no other triggers that may demand both?

MR. GEDDES: There are triggers that certainly trigger both, yes. If you go back to that slide, that's the next column over.

That potential CCF column means we've recognized that the triggering conditions can be concurrent.

MR. TOROK: Now, coming back to the sensor, if the senor fails and that fails the channel, and you have annunciation in the control room that the sensor failed, so you know the sensor failed, then the likelihood of having multiple sensor failures that causes this common-cause failure in the software is very low. So that one becomes unimportant.

If, on the other hand, the failure in the digital system is such that nobody is looking for it,

and nobody knows it happens, and then it happened, and then the next channel fails, now you really do have a potential common-cause failure. We saw that as well, where, for example, one power supply failed, and three weeks later a back-up power supply failed, and nobody knew in between. Well, that one really was a potential common-cause failure because of the way the system was designed.

MR. SIEBER: Let me ask this question: if you are examining operating experience by looking at LERs, the only ones you will find are in the far right column in the LER, right?

MR. GEDDES: I beg to differ. We found several that were in those third and fourth columns.

MEMBER BLEY: That's interesting.

MR. SIEBER: But no event. How would you find them in the LER?

MR. GEDDES: Here's why. Let me take a shot at it.

There is events on a core potential calculator system where they discovered a software defect. In fact, the vendor discovered it and told the plant. The plant reported it. The plant --

MR. SIEBER: As what, a Part 21?

MR. GEDDES: Well, they shut down --

MR. SIEBER: It's not an event.

MR. GEDDES: They shut down their plant. Okay?

MR. SIEBER: Well, it's the shutdown --

MEMBER BLEY: Generated the report.

MR. GEDDES: Right. They entered a 303 action statement that said, "We're inoperable. We need to shut down."

In hindsight, they realized, after the dust cleared, that they were operable but degraded. They had the presence of a common defect. They had exactly in that middle column. Okay?

They incorrectly assumed that the presence of a common defect meant all four channels were inoperable. Now I know this is a tech spec --

MEMBER BROWN: But why wasn't that reasonable in the context of you don't figure that out immediately?

MR. GEDDES: Because Generic Letter 91-18 allows the idea that you can be operable but degraded. In this case, it would have taken four concurrent sensor failures to render all four channels inoperable because it was a sensor failure that triggered the fault.

MEMBER BROWN: Yes, I know. I understand

that.

MR. SIEBER: On the other hand, that is a specific example, I believe -- and you can correct my belief, if you would like -- that not all of these first four columns will end up as an LER report.

MR. GEDDES: That's true, and we do have some INPO OE reports where they wouldn't pass the threshold for reportability to the NRC, but they did come out in INPO OE reports.

MR. SIEBER: Now are the INPO operating experience reports a major part of your database?

MR. GEDDES: Yes. It's half.

MR. SIEBER: Okay. Because LERs by themselves don't tell the entire story.

MR. TOROK: Right. Actually, I think we started out characterizing this center column, where you had a common defect, as a potential common-cause failure. But when we got into discussing them with the group, somebody said, "Wait a second. The trigger is outside there. You can't really make that into a common-cause failure, no matter what you do."

MR. GEDDES: And that's why I did those text boxes in those screenshots. Okay? And I started off saying, if there's a common defect, it is a potential CCF, period. My peers said sometimes they

are; sometimes they are not; it depends on how they are triggered.

I went back and I started reclassifying. There must have been 40 or 50 events that I combined those third and fourth columns. I started to split them out. I'm not sure I went back and fixed all the commentary. Okay? That is one of the issues.

MEMBER BLEY: That's kind of where I am hanging on this. If we start at the left, we've got no problem. Then we get a real single failure. Then we have the no common cause. Does that always mean that we have had a common defect, if something is labeled no common cause?

MR. GEDDES: That means the event report reported the presence of a common defect. They didn't actually have a failure, but it could have failed in the form of a single failure.

MEMBER BLEY: So somebody who wants to go through your data and give maybe more thought to these, are there other triggers being thought about, can find them because they are labeled no common cause failure?

MR. GEDDES: Well, in the potential CCF, that dotted red line means we have discovered and classified the triggers.

MEMBER BLEY: Okay, and it doesn't mean this event actually had a trigger. It means you found there was a potential trigger?

MR. GEDDES: It could have. It could have triggered. In the far righthand column -- the only two places where a defect was triggered, either in a single failure or a common-cause failure, are the second and the fifth, where there's those red boxes.

MEMBER BLEY: But the fourth box, with the dotted lines, does that mean that the particular event you were evaluating actually had a trigger present or that you were able to divine a trigger that could have actuated this event?

MR. GEDDES: Well, it would be like a Part 21 report that said, "We found a software problem, and it could result in inoperable -- it could result in common-cause" --

MEMBER BLEY: I'm not saying my question right.

MR. GEDDES: Okay.

MR. TOROK: We found concurrent -- we found the possibility of concurrent triggers, is the answer, I think. That is why we called it, we put it in that fourth column.

MEMBER BLEY: You had an event. You

looked through it. You identified that this particular event, indeed, had a common defect.

MR. TOROK: Yes.

MEMBER BLEY: And you identified, whether or not there was a trigger, you identified that there was a possibility of concurrency?

MR. GEDDES: Correct. Of concurrent triggers. And that is the difference between those two.

MEMBER BLEY: And if you identified triggers, but they are not concurrent, you gave an explanation of why they wouldn't be concurrent? I don't remember.

MR. GEDDES: In our discussions we sure did. I don't know how well that is documented.

MEMBER BLEY: Anyway, they can be found?

MR. GEDDES: Yes.

MEMBER BLEY: Anywhere there's a concurrent defect, they can be found, and anybody who wants to can think about those as hard as they want?

MR. GEDDES: Yes.

MR. TOROK: You could go back and revisit all the events that were called common defect and ask yourself that question: are there concurrent triggers or is there a possibility of that?

MEMBER BLEY: I don't want to hang up any more on this. I wanted to understand what was there and what somebody could do with it.

MEMBER BROWN: I want to ask one other question in terms of the common defect, a trigger, not the common defect, but the trigger.

This was a digital system, correct? I mean we are still just working with digital systems?

MR. GEDDES: Yes.

MEMBER BROWN: Okay. Typically, you have a detector, which is an analog device.

MR. GEDDES: A sensor.

MEMBER BROWN: A sensor. Yes, not a detector. I'm sorry. Other program, old program.

Then you normally have, typically have some signal condition of some kind. That is an analog function.

Then it goes in an A to D converter. Now you have data being put out into some buffering system, memory, you know, some points where it can come through and sample those.

So, when you talk about the sensor defect, trigger -- excuse me -- trigger, do you know the nature of that trigger? Was that reported by the vendor?

What I am looking for was that, normally, the sensor runs through a range. Take a temperature detector, you know, from 100 degrees to 600 degrees, and it goes through the conditioning. It outputs to the A to D converter, and you have a range of bits and bytes that then can be sampled and picked up.

Was this a bits and bytes?

MR. GEDDES: No. No, it goes to the failure modes of the sensor itself in that example. Fail high; fail low.

MEMBER BROWN: Yes, but that means that the bits and bytes go to either 600, using my example, or 100.

MR. GEDDES: Right, and the software defect was that the resulting output of the system will not meet the requirements.

MEMBER BROWN: So did they do range checking?

MR. GEDDES: In that particular case, they did, but it was implemented incorrectly. There was a logic issue in a software build that was delivered after the upgrade, and the vendor reported saying, "That software build incorrectly codifies this range-checking algorithm and will result in a channel failure if you get a sensor failure."

MEMBER BROWN: Because a particular mode of a sensor failure ends up with an A to D conversion, a set of bits and bytes --

MR. GEDDES: That is more in the logic, the design itself --

MEMBER BROWN: Yes.

MR. GEDDES: -- of the algorithms.

MEMBER BROWN: Yes, but, fundamentally, you get to the point, the initial setup is you are either low or you are high. Theoretically, if you are in the middle, everything works okay. It either failed low or high in this circumstance.

MR. GEDDES: In this case, yes. There are fail as is failure modes that we consider as well.

MEMBER BROWN: Well, it could fail as is, but if it is in the range, then it looks like a piece of data --

MR. GEDDES: Then we do channel checks to compare.

MEMBER BROWN: I don't want to even start on that. Okay?

MR. GEDDES: Okay. I appreciate that.

MEMBER BROWN: I don't care what that channel does, just as long as the other ones are still in place.

MR. GEDDES: Right.

MEMBER BROWN: All right, I got the picture. Go ahead.

MR. GEDDES: So anymore questions on this taxonomy before we go further?

MR. HECHT: If there are multiple divisions or channels that are dependent on a single datapoint, that wouldn't be true in the safety case, but it might be true in the control case, is that correct?

MR. GEDDES: There are cases in control system implementations that we saw that were dependent on single shared resources, like power supplies, sensors.

MR. HECHT: Right.

MR. GEDDES: Of course, that could lead to a master-slave failure mode, a concurrent failure mode.

MR. HECHT: So a single failure in the situation like that would actually be called a concurrent trigger?

MR. GEDDES: It would be a common defect and a concurrent trigger, yes.

MR. HECHT: Okay.

MR. TOROK: So the definitions are a

little different for non-safety. That is also, what you pointed out was one of the reasons why it is difficult to combine the safety data with the non-safety data, because you get into those kinds of things, where it can happen in non-safety, but safety doesn't have that.

MEMBER BROWN: It is interesting. Thank you for what you just said. You triggered another thought.

CHAIR APOSTOLAKIS: Too many triggers here.

(Laughter.)

MEMBER BROWN: I am sorry. I get wrapped around an axle every now and then.

When I think about this triggering issue, here's a set of data coming in from the sensor. It gets converted. The output of that converter is then sampled, blah, blah. That is where, whether it is high, low, whatever that range thing is.

If that data -- we talked about channel-to-channel stuff -- if that data then is in a design where they share the data from channel to channel, I have now -- you just told me it wouldn't work in that channel. It is zeroed out.

I have compromised my entire protection

system when I do that, unless you have some really sophisticated algorithm, artificial intelligence, that says, oh, when I'm examining all this data to decide which one I am going to use in this non-feedback control system -- I'm thinking reactor protection system, when all I want to do is shut something down when I am not in the right place.

In other words, we are going to be smarter and share this type of data between the things. It is a potential problem when you start compromising that independence from channel to channel.

MR. GEDDES: ISG 4 has very specific criteria.

MEMBER BROWN: I have read those. They are abhorrent in some circumstances.

That's a good word. I thought you would like that word, George.

(Laughter.)

MR. SIEBER: But how do you like them?

(Laughter.)

MEMBER BROWN: I wanted to make it clear. Okay? Because ISG 4 does talk about the shared data issue, and it creates some real problems.

MR. TOROK: That is a fair question: what are you doing to avoid this problem?

MEMBER BROWN: Yes, well, in a feedback control system sharing data, like I said, when you want fault tolerance, in a protection system where you want independent channels to be really independent, then you share that data. If it is in the name of fault tolerance in those channels in the protection system, that is a bogus thought process. You don't care. Okay?

You were always saying one channel doesn't work, and you assume this one doesn't work and I've got my other ones remaining.

I just have to get a few thoughts out here philosophically. I am ready to go on, George.

MR. TOROK: Yes, and we leaped into ISG 4 there a little bit.

Are we okay on this guy?

MR. GEDDES: Mr. Hecht, did we answer your question?

MEMBER BROWN: I interrupted him. I'm sorry.

MR. GEDDES: Did Mr. Brown answer your question?

MR. HECHT: I think you answered it very well when you said, basically, concurrent trigger does not necessarily mean that there are multiple events

that have to happen or multiple initiating events.

MR. GEDDES: That is true. It depends on the design.

CHAIR APOSTOLAKIS: But we don't have, Dennis, we don't have a similar situation in hardware, similar to the third column, do we? I mean we never distinguish. I mean we say the demand is demand.

MEMBER BLEY: We almost never find those.

CHAIR APOSTOLAKIS: Yes.

MEMBER BLEY: I mean they exist out there. You could have a manufacturing flaw in a bearing or something.

CHAIR APOSTOLAKIS: No, but the flaw is, but we don't make a distinction between triggers. We say there is a demand for high-pressure injection; all drains have demanded.

We don't say, oh, gee, in this situation --

MEMBER BLEY: People have played with that. We haven't developed that well enough.

CHAIR APOSTOLAKIS: Yes, I don't think we have. Yes.

MEMBER BLEY: It would be a good thing, though, because there are some things classified as common-cause failure of mechanical equipment that

probably shouldn't be because of the same --

CHAIR APOSTOLAKIS: The same thing?

MEMBER BLEY: -- kind of thing. You can't get to the trigger.

CHAIR APOSTOLAKIS: But I think this is more appropriate here, though.

MEMBER BROWN: No, but your point is valid. This is a classic example of, can you really do software testing, V&V, prior to your execution, where you are going to catch all of these latent defects?

There will be latent defects. You will not catch them all. So you have to depend on some armor belt, independence, which is a big one, a very big one, which --

MEMBER BLEY: Now we are getting outside of --

MEMBER BROWN: And we are getting way outside.

MEMBER BLEY: -- what's even tested.

MEMBER BROWN: Exactly right. So anybody that thinks we can have enough testing regimens where we are going to find all these, they are going to be there. Therefore, you have to set up your really critical protection and safeguards channels with the

only protection you've got, which is independence. You don't want to compromise channel to channel.

MR. GEDDES: Well, I would like to add that, based on an earlier comment, the purpose of the research was to answer the basic question that I think was raised by the ACRS: what does the OE tell us? Okay? This is what it tells us.

CHAIR APOSTOLAKIS: Yes.

MEMBER BROWN: It is good. I mean this is great to see somebody surveying this stuff and laying it on the table.

MR. GEDDES: Mr. Brown, to piggyback on your comments, where we try to protect against CCF, we go after to prevent or reduce the presence of common defects, and we go after reducing or preventing concurrent triggers. Independence helps immensely in the context of triggering. Okay?

If we can limit a trigger through the use of defensive measures, for example, to a single channel, that is a means to combat CCF. Of course, we are always going after common defects.

MEMBER BROWN: Oh, of course you do, but you're just never going to find them all. That is the problem.

MR. GEDDES: Right.

MEMBER BROWN: We've got to recognize that.

MR. GEDDES: We think it is a two-pronged approach here.

MEMBER BROWN: I agree.

MR. GEDDES: And the OE helps us find where they are.

MEMBER BROWN: I agree. It is very, very useful.

CHAIR APOSTOLAKIS: I mean, we keep talking about failures and defects that were triggered, and so on. What was the actual failure mode?

I remember Myron here brought to us a classification from some other industry, which we put in our letter, "hung" and "delayed".

MR. TOROK: Well, we will get to that.

CHAIR APOSTOLAKIS: Are they doing that?

MR. TOROK: We will get to that in a little while.

MR. GEDDES: In this construct, we talk about system-level failure modes, loss of function at the system level.

CHAIR APOSTOLAKIS: Even there, I mean what does it mean to lose function? If you come back,

that's fine.

MR. TOROK: But, for the purposes of this study, this OE study, what we cared about the most were the last two columns on this chart in terms of, what's the operating experience in regard to actual common-cause failures and potential common-cause failures, which means there is a common defect and the possibility of concurrent triggers, those last two columns. That is really what we were trying to isolate out of all these 322 events. We are trying to find those. Okay?

MEMBER BLEY: That middle column still is hanging up for me a little bit. I don't know how well you did this or how thoroughly. But I think you did the looking for the concurrent triggers in a collegial discussion kind of arrangement.

The systematic approach used in systems analysis might uncover things that you don't uncover in that process. I'm not saying you should have done that. I'm just saying somebody else might want to take a more systematic look at those things in the third column and see if, for particular systems, if there's a potential problem hanging here that we didn't identify that way.

MR. GEDDES: There is particularly useful

information in the reports in terms of the corrective actions. If there is a formal root cause, we use INPO methods for really getting down to the root cause. The idea of the corrective action of reconcurrence is, what one thing do we have to do so this never happens again?

Where we get into those middle column events, the root causes and the corrective actions can be very revealing.

CHAIR APOSTOLAKIS: Okay.

MEMBER BROWN: I won't dispute that.

MR. TOROK: May we?

CHAIR APOSTOLAKIS: Yes.

MR. TOROK: Bruce?

MR. GEDDES: Okay. Somebody asked the question, what do we mean by failure modes? We found that there were no actual CCF hard failures that completely disabled the safety function at the system level. Okay? That is one of the first findings.

We found actual and potential CCF events were dominated by non-software issues, life cycle management and human performance errors.

MEMBER BLEY: I just want to hang on that first one.

MR. GEDDES: Okay.

MEMBER BLEY: Most of these systems that are out in the field aren't in places that would completely disable a safety function. But go ahead.

CHAIR APOSTOLAKIS: What do you mean by function?

MEMBER BLEY: And safety function? Do we mean --

MEMBER BROWN: Well, I'm not aware of any --

MEMBER BLEY: Is feedwater a safety function? Probably not.

MEMBER BROWN: No, it's an event-causing function if something fails. But reactor protection systems and safeguard systems, are there any digital INC ones out in the U.S. plants today?

MR. GEDDES: Yes.

MEMBER BROWN: I presume there are. I just don't know what they are.

MEMBER BLEY: How many?

CHAIR APOSTOLAKIS: So a safety function, I mean injecting water, that kind of thing?

MR. TOROK: Right. Exactly.

MR. GEDDES: On demand.

CHAIR APOSTOLAKIS: The operating experience is very limited then, right? Charlie, I

think the operating experience --

MEMBER BLEY: And Bruce just added on demand. If there was no demand, then --

CHAIR APOSTOLAKIS: See, that is too strong a statement. I mean you have very limited experience because we don't favor widespread use of these systems. Now you say on demand. In other words, you are waiting for a LOCA?

MR. TOROK: Well, yes, in the data we looked at -- well, that's a good point because this is basically saying, look, in 1E systems we didn't see any actual common-cause failures. You shouldn't expect to, right, because they are not called upon to act very often, and the systems are, by design, very robust. So that shouldn't surprise anybody. Right?

CHAIR APOSTOLAKIS: But surely you don't mean that I have to have a LOCA in order to say I had a common-cause failure?

MEMBER BLEY: You just triggered something for me.

CHAIR APOSTOLAKIS: I did?

MEMBER BLEY: The fourth column then, the concurrent defect and concurrent triggers, at that point, if you had a demand, you failed?

MR. TOROK: That's right.

MEMBER BLEY: So that fourth column, the system is dead. You just don't have a demand. That is why it is not showing up as a common-cause failure.

MR. TOROK: That is the difference between the third and the fourth really.

CHAIR APOSTOLAKIS: But I think, when we study software, I mean it is like if we studied hardware. I have a system. I want to know whether my system will fail in a hypothetical demand due to software failures. That is really the focus.

The fact that I haven't had an accident sequence where that would have played a part, yes, I mean, gee, when we do hardware analysis for the various injection systems, and so on, we assume that there has been a LOCA, and then we do the analysis. We never say, but the high-pressure injection system never failed because we never had a LOCA.

MEMBER BROWN: But the latent defects issue, you've got the latent thing in two places. Okay, it's there.

MR. TOROK: That's the fourth column.

MEMBER BROWN: But here, in the fourth column, you have to assume that two sensors both produced -- in other words, you're going to have a double failure.

CHAIR APOSTOLAKIS: So the last one is a real one?

MEMBER BROWN: The last one is, yes, that I'm not quite sure I understand.

MEMBER BLEY: On demand, whatever this is starting didn't start.

MEMBER BROWN: Yes. Column 4, if you had a LOCA, you can't tell whether that one would actually not respond or not.

MR. TOROK: No, it would not respond in the fourth column if we had the LOCA.

CHAIR APOSTOLAKIS: It would not. It would not have.

MEMBER BROWN: You have to have two sensors though. You have to have two sensors.

MEMBER BLEY: No, no, no. Only for that one example.

MR. GEDDES: That's just one example.

MR. TOROK: In general, in our taxonomy, if you are in the fourth column and you have the trigger, it is concurrent triggers.

MEMBER BROWN: Okay. So you're saying that the other sensor is going to have that problem somehow.

MR. TOROK: The other trigger, whatever it

is --

MEMBER BROWN: And it is going to occur along with it?

MR. TOROK: Yes, that event can have concurrent --

MEMBER BROWN: So they're saying you're getting concurrent triggers.

MR. TOROK: Yes.

CHAIR APOSTOLAKIS: So the great interest is the last two columns on the right?

MR. TOROK: That's right. Those are the two you care about, and you care about both of them. So this notion that you have never had an actual common-cause failure is not very reassuring because that fourth column still counts.

CHAIR APOSTOLAKIS: You care about them from the point of view of the consequences, failure analysis. But it seems to me that the other columns, especially the third one, would be useful in terms of understanding what kinds of failures --

MR. TOROK: Absolutely.

CHAIR APOSTOLAKIS: -- find their way into software. So I think they contain useful information.

MR. TOROK: That's interesting you should mention that. This is not really part of this study.

However, we did take the results of that type and use them in another EPRI project where we produced guidance along those lines and training materials along those lines to capture those --

MEMBER BLEY: And you would want to use those if you were trying to model.

MR. TOROK: Sure.

CHAIR APOSTOLAKIS: Sure. We started moving back --

MEMBER BROWN: I still walk down the path on that fourth column, Dennis. Somehow I've got to postulate a sensor design that generates a concurrent -- where the next sensor is going to produce the same thing in a concurrent manner.

I understand the philosophical argument in which you generate it.

MR. TOROK: Forget about sensors. We will give you an example.

MEMBER BROWN: It's external. If that's shared data, then I agree with you. One can do it and trash your whole system. I agree with you.

MR. TOROK: We will show you an actual example that has concurrent triggers. Okay? I think it is like the next slide almost.

MEMBER BLEY: You may not want them, but

there are places where one sensor feeds more than one --

MEMBER BROWN: I understand that. We've already been told of that.

MR. GEDDES: You will notice in the second bullet on this slide we lump actual and potential together in the same idea. We don't exclude potential CCFs.

The third bullet, we found that current methods suggest that they are effective in keeping software a minor contributor. We are not proposing that software -- you know, our interest in software quality assurance and the way we manage software should set the concept of the CCF aside, we don't mean that at all.

But, as Ray mentioned earlier, we want to investigate those methods even further and leverage them, so that we can keep this trend low, like we have seen.

MEMBER BLEY: Without some comparison to the other contributors to common-cause failure, you can't make that statement. So where did you make that comparison?

MR. GEDDES: The next slide.

MR. HECHT: I have a question on the

previous slide. I'm sorry.

With respect to you use the words "life cycle management". Do you mean configuration management or do you mean something else?

MR. GEDDES: All of it. Requirements analysis, V&V, configuration control.

MR. HECHT: So it is the entire software development and implementation process.

MR. GEDDES: We call out software as design or logic errors. For example, the requirements were complete and correct, but the software itself incorrectly implemented logic that did not meet the requirement.

MR. HECHT: Okay. So that is also life cycle management.

MR. GEDDES: No, that would be a software issue. Everything else is a life cycle issue.

MR. TOROK: Like setting the setpoints for the system or calculating the setpoints and implementing those or parameters that the system needs to operate --

MR. GEDDES: The requirements there in our construct here would be considered a life cycle issue, not a software logic issue.

MR. HECHT: A requirements error would be

considered a life cycle issue. A setpoint would be considered a life cycle issue. What if somebody loaded an incorrect version of the software?

MR. GEDDES: That would be a life cycle issue. Let's say they pulled an out-of-date version off the shelf.

MR. HECHT: Okay.

MR. GEDDES: That's a human error.

MR. HECHT: I guess that would be the software requirements, not the system requirements. But everything between the software requirements, specification, and what about tests?

MR. GEDDES: If there is a testing error?

MR. HECHT: Yes.

MR. GEDDES: Testing errors usually aren't a root cause. Inadequate testing might be a contributing cause to failure to discover software logic defect. Okay? But that would be considered a life cycle issue.

MR. HECHT: Okay. So it is really design and coding errors that are excluded from life cycle management?

MR. GEDDES: Yes, that is what we mean.

MR. TOROK: To a large extent, it is the processes that are used right now for configuration

management of existing systems. How do you control the configuration and the setpoints and the parameters on the analog systems? That is all in place now. It is done under Appendix B programs, and so on. There are processes for that. Right? They are not peculiar to digital or software. So we tried to separate that.

MR. HECHT: Okay, got it.

CHAIR APOSTOLAKIS: Regarding this line, what is the OE telling us? In reading the report, I noticed that you don't miss an opportunity to say that diversity is not helpful.

For example, "This event also shows why platform diversity is not always effective."

MEMBER BROWN: Which report are you in?

CHAIR APOSTOLAKIS: The operating experience.

MEMBER BROWN: Oh, okay, the OE? All right.

CHAIR APOSTOLAKIS: Then I suspect I know why you are saying that, but on page 55, "The majority, 18 of 27, common defect events in 1E systems resulted in subsystem or channel effects, leaving the balance of the system unaffected and available to perform its overall safety function by other means, using functional or signal diversity."

Well, that's actually positive.

Are you saying that diversity is not necessarily a good idea?

MR. TOROK: Not at all. All we are saying is, in the events we looked at, there were some really good examples of where functional diversity and signal diversity were obviously helping. There were no examples that we could see, that we saw in the events we looked at, where platform diversity was advantageous. That's all.

So certain types of diversity are certainly very valuable and you don't want to give them up. No doubt about it.

CHAIR APOSTOLAKIS: So I guess the message we are getting here is that, when people talk about D3, we should just try to apply it blindly. I mean there are situations where the diversity part is useful, but in other situations it might not be.

MR. TOROK: Exactly.

MR. GEDDES: There's multiple forms of diversity, and we have to be careful about which forms we apply.

MR. TOROK: Right. Diversity necessarily adds complexity, but it doesn't necessarily add safety. So you want to be judicious about that.

CHAIR APOSTOLAKIS: The report we reviewed, though, from Oak Ridge, diversity was really the No. 1 --

MR. GEDDES: Equipment diversity was No. 1.

CHAIR APOSTOLAKIS: Huh?

MR. GEDDES: That report puts heavy emphasis on equipment diversity.

MEMBER BLEY: Did you have enough source information to really conclude that platform diversity --

MEMBER BROWN: There was no bridge report. I am not sure they know --

MR. TOROK: We didn't say platform diversity wasn't valuable. We just say we didn't see any cases where it was.

MEMBER BROWN: But it might be because you didn't see many places where there was platform diversity.

MR. TOROK: That's right. That might be. But one of the things we asked ourselves --

CHAIR APOSTOLAKIS: Maybe I'm wrong.

MEMBER BROWN: I don't remember that. I'll have to go back and look.

CHAIR APOSTOLAKIS: Okay. So you are not

against diversity?

MR. TOROK: No.

MEMBER BLEY: We just had a little conversation while yours was going on.

I had asked, did they really have enough places where they have seen platform diversity to draw the conclusion that it wasn't helpful? I think the answer was it hadn't been helpful in the events they looked at.

CHAIR APOSTOLAKIS: There was one.

MR. TOROK: In looking at each event, what we asked ourselves, once we thought we understood the event, was, what would have been helpful here in terms of defensive measures or in terms of diversity attributes?

MR. GEDDES: Different forms of diversity attributes.

MR. TOROK: There's some good hints in the writeups in terms of what the corrective actions were. Right? So we always ask ourselves that question.

That is why in some cases functional diversity and signal diversity jumped out at us. We said, wow, these guys saved the day here. Right?

But there were none that we saw where platform diversity looked like an advantage; that's

all.

MEMBER BLEY: Of course, it may have some of the others embedded in them, as you go from one platform to the --

MR. TOROK: If you had gotten into a situation where you were seeing a lot of failures coming from operating systems and platforms, or something like that, then you would have said, wow, platform diversity would have helped here.

MEMBER BLEY: Okay, fair enough.

MR. TOROK: But we didn't see that, that's all, in the stuff we looked at.

MEMBER BROWN: From a platform -- go ahead.

MR. HECHT: Okay. Well, I'm on slide 17, which I looked at earlier. There are certain things, such as processor malfunction, EMI, and I recall it's not reported.

MR. GEDDES: Give us a chance to get caught up here.

MR. HECHT: I'm sorry. I'm just looking at those. There are some mechanisms which seems to imply that maybe it would be.

MR. GEDDES: I'm sorry. Say it again?

MR. HECHT: So processor malfunction,

there was one incident like that. Root cause not reported. In other words, they didn't know what it was. Or EMI.

CHAIR APOSTOLAKIS: You are on which slide?

MR. HECHT: Seventeen. I'm just pointing out that --

CHAIR APOSTOLAKIS: On slide 17?

MR. TOROK: We would have to go back and look at those individual events and go over them in detail really to respond to that probability, but that is a good question.

MR. GEDDES: Now, for a lot of these, these are human error and life cycle management. EMI, for example, we use qualification methods, and an effective qualification program could equally affect two different platforms, if the tests were inadequate or the specifications are incorrect.

MR. HECHT: Well, except that at the time of the actual incident, it could be that the designed diversity, and it's not even software there -- it is probably a hardware design diversity issue.

MR. TOROK: Now, of course, in principle, your qualification program for EMI is supposed to address the issue of EMI, right?

MR. HECHT: Well, that's almost like saying, in principle, your software test is --

MR. TOROK: Well, if you talk about adequate assurance of how to get protection against various things, the EMI one is handled through qualification really. But that doesn't, what you are saying, that doesn't negate your comment here at all. It is just an observation that, for a number of causes, they are addressed by other means, through normal qualification processes, and so on, right now, and that is considered adequate.

Shall we go back?

MR. HECHT: Okay. I'm sorry.

MR. TOROK: Okay, so we are on to this one now, I think.

CHAIR APOSTOLAKIS: Which is No. 13?

MR. TOROK: No. 13, 1E common defect events.

MR. GEDDES: Okay. This construct on the lefthand side you have seen before. We break down the events. We start at 322. Forty-nine report something on a 1E system. Out of those, we see 27 cases where a common defect was reported, and focused particular interest on that.

The single defects are interesting, but

our primary focus was on, what about these common defects?

The way we have defined our terms, four of them are software- or logic-related. Okay? Twenty-three are not.

Mr. Hecht was peeking ahead, and we just looked at some of those. Okay?

But, in this case, we used that chart taxonomy on the righthand side to break it down where we had six of these potential CCFs out of the 27 and no actual CCFs, the way we have defined those terms. Okay?

Out of those six potential CCFs, one of them was software-related and five were not. The six potential CCFs -- you can see the balance of the 27 common defects are down below.

We say, for example, 10 single failures, but that is where the triggering condition would result in that middle column. It is a common defect, but the triggering condition means it can only manifest itself in one channel at a time, not concurrently. And two of them are due to software. We are not hiding those software events, but they don't result in a potential or an actual CCF.

MEMBER BLEY: I am going to take you back

to your one software potential common cause. Well, actually, I'm taking you to your five other ones.

Because of the way you define software, though, if something in the process ended up corrupting the software, that would not be catalogued as a software common-cause failure? That is what you told me earlier.

MR. TOROK: What do you mean? What is an example of something in the process?

MEMBER BLEY: Data comes in and a register overflows and somehow screws up the code.

MR. GEDDES: It depends on the mechanism. If it's caused by an operating system defect, then we will would classify this as software defect.

MEMBER BLEY: But if it is caused by some other situation?

MR. GEDDES: Like BMI, for example, or --

MEMBER BLEY: A cosmic ray.

MR. GEDDES: -- the ubiquitous cosmic ray? We would not call that a software defect.

MEMBER BLEY: Or input data outside the range of where it was tested?

MR. GEDDES: Right. We wouldn't call that a software defect necessarily.

MR. TOROK: Well, I don't know about that.

We would discuss that one.

MR. GEDDES: Yes.

MEMBER BLEY: You didn't do it.

MR. TOROK: No, no, no. Thuy would jump in. He would jump all over us on that one.

MEMBER BLEY: Yes? He would?

MR. GEDDES: Get the microphone, Thuy.

MR. NGUYEN: Thuy from EDF.

I would say, in the case of an incorrect input that would cause the software to crash, that's for me a software issue because the first thing you have to do, when you have inputs, is to verify that your inputs are in the correct range.

MEMBER BLEY: I like what you are saying, but earlier we were told that almost anything -- that corruption of the software wouldn't lead to a classification of software.

MR. GEDDES: It depends on the error. If the requirement is, the range is, I don't know, zero to 1,000 pounds, and the real range should have been zero to 1200 pounds or zero to 800 pounds --

MEMBER BLEY: It is not that it should have been. It is somehow the real world took you outside of --

MR. GEDDES: Okay, that's a better way to

put it.

MR. NGUYEN: In the case you are speaking about, it is a software issue. However, the trigger is a random trigger. So the fact that the input is incorrect does not necessarily affect the four channels at the same time.

For example, in our case, the way we enter new parameters, new parameter values, in a safety-redundant system is we do it one channel after the other. So we verify on one channel, wait for some time, sometimes 24 hours, for example, and then do it on another channel. So, if the fact that we enter an incorrect value causes a problem and causes the digital system to crash, it will affect only that channel.

But, still, it is a software problem. The software is not supposed to crash, whatever the input values.

MR. TOROK: Now this gets into a discussion of what's adequate in terms of defensive measures, too, because I think Thuy would say any software system worth its salt is going to know what to do with any possible input it can see. You know, there's an anticipated range where this thing goes, but the software should know what to do if the input

goes outside that range.

Of course, you can make software do that.

MEMBER BLEY: But the other side is it's not tested.

MR. TOROK: That's what data validation is about.

MR. NGUYEN: And in fact, it is not up to the software engineer to decide what to do when you get incongruent values that are out of range. That must be part of the system requirements specification.

MR. GEDDES: Exactly. And if the requirement specification doesn't adequately describe the real world, that is a requirements problem.

MR. HECHT: Which is outside of the software.

MR. GEDDES: If it is outside of our definition --

MEMBER BROWN: If it is outside of your life cycle management, which is a comment you made earlier.

MEMBER BLEY: But the problem I am hanging on is there have been, not in our nuclear systems, but there have been cases in some power control systems and in some medical places where this kind of problem has occurred. If it occurred in an analog system, you

would get something wrong, but once the data cleared, everything would be right. But if you get in a spot that somehow that leads to a corruption, then you never recover with these kinds of systems.

If we are putting those events in another bin, then the kind of problems that have led to blackouts in the Northeast and to some deaths in the medical business are getting pushed out of our look for common-cause failures.

MR. TOROK: Come back to this notion of requirements for a minute though. When we looked at the events, if there was an event where there was a requirement specification omission or error, or something like that, that led to the event, right, one of the questions we would ask ourselves is, suppose this system had been implemented in analog technology. Would it have had the same problem? Right? Sometimes the answer is, well, yes, because you can't tell the difference in software written on bad requirements, in which case we have said then it is not really a software event.

MEMBER BLEY: Well, it is important, and we want to understand that.

MR. TOROK: However, but also, if the thing had been implemented in such a way that the

digital system added a lot of functionality that the analog system didn't have, simply because it couldn't --

MEMBER BLEY: Or had a failure mode.

MR. TOROK: Now we are talking about definitely a software problem. Right?

MEMBER BLEY: Okay. If that's true, I'm happier.

MR. TOROK: We tried to make that distinction.

MEMBER BLEY: We haven't had those events yet. So that's our concern about it.

MR. TOROK: We had extensive discussions about that. I'm not saying we got them all right, but we tried to do that. Okay?

MR. HECHT: May I point out that there are some aspects of digital systems which are quite relevant in this regard? The whole example of the fact that you have a cycle where you sample and then evaluate and then put out, I mean you don't have that in an analog system, of course.

MR. TOROK: That's right.

MR. HECHT: It's all continuous.

Another example would be the Nyquist frequency. A third example would be D to A issues

which you don't have in an analog system.

MR. TOROK: Right.

MR. HECHT: So there's a whole class of requirements that you might consider to be excluded, but which, in fact, are indirectly caused by the nature of the --

MR. TOROK: I don't think we would exclude the ones you just said.

CHAIR APOSTOLAKIS: You don't think what?

MR. TOROK: We would not exclude the ones Myron just said. Those would still be in there.

If the problem results from a behavior that is peculiar to digital technology, we are going to call it a software problem, right? Regardless of whether it comes from the software itself or the digital system architecture or something like that, we are going to call it a software problem.

MEMBER BROWN: But a D to A converter can fail independent of the software data coming into it. I mean it is a device. So, independent of what the software is doing, it can have a failure mode, and you have to account for that --

MR. TOROK: Yes.

MEMBER BROWN: -- in your system design. That is one of your single failures you have to deal

with.

Go ahead. I'm sorry.

MR. HECHT: But, in that case, the resolution might not have been right. The head room might not have been right.

MEMBER BROWN: I agree. I mean there's a lot of different things. If you exceed a range on a converter, you can have problems if the converter is not one that will accept that very well. You have to do something else to ensure that it doesn't exceed its application range. But the older ones had that problem. The newer ones don't necessarily have that problem these days.

MR. AUSTIN: Rob Austin with EPRI.

I think we are all on the same page here, but what we don't want to say is that, when we put the stuff through the sieve, the only ones that we are concerned about as an industry are the common-cause failures and software that pop out at the bottom.

We are concerned about all of them, and that is one of the major learnings, is that there are a lot of other ways besides software that you can step into it with these systems. We are looking into that.

For example, we are taking the learnings in this OE database and we are starting a project on

better maintenance for digital systems, because if you go and look at some of the causes, you will see inadequate PMs, and that is unacceptable from an industry point of view.

Another example of this is the failure analysis that we talked about before. We are not going to focus just upon the software common-cause failure, but the whole range of failures.

So I don't want to give the impression that -- I think we're all in agreement, but I just want to say that it is not just -- you don't want to put so much focus on the software common-cause failures that we forget about other stuff which is equally a source of problems.

It is also this whole definition shows, when I said earlier the importance of at least being in agreement on what software is, and it may be a case where we can't please all the people all the time, but at least there is a common agreement. It does become elastic sometimes.

The definition of software becomes even tougher when we get into FPGAs and other type devices that are coming down the pike.

Thank you.

MR. TOROK: Are we done with something?

CHAIR APOSTOLAKIS: Well, what kind of failure was this subsystem actual failure? What did it do? I mean, did we respond to demand or what?

MR. GEDDES: Which one are you looking at?

CHAIR APOSTOLAKIS: The very last one. It says, one subsystem actual common-cause failure. A setpoint issue.

MR. GEDDES: Right. That was a case where there was a reactor trip lightning strike on the main transformer. The plant reacted to a loss of the main transformer. There was a time delay in the reactor protection system. I think it was a core protecting calculator instance. A subpoint for detecting rod motion was incorrect. The actual parameter itself didn't account for the real-world case of how far a rod can slip in a certain amount of time. And then sort of a second time delay, there's a 16-second time delay related to the way the rods are supposed to behave under certain transient conditions.

CHAIR APOSTOLAKIS: There was a delay.

MR. GEDDES: I'm sorry?

CHAIR APOSTOLAKIS: It was a delay.

MEMBER BROWN: You mean the rod started to drop and then recovered because of the lightning strike?

MR. GEDDES: The rod slipped a little bit. The reactor was supposed to trip the first time, but it inserted a second time delay, a 16-second time delay, and eventually tripped 16 seconds after it was supposed to for this particular transient. We call that, that's a failure on demand, where there was a plant condition where the reactor should have tripped, and it was delayed by another time delay that it wasn't supposed to do that. Okay?

MEMBER BROWN: Yes, I know, I understand.

CHAIR APOSTOLAKIS: Is that event 222, right? That's event 222? I believe it is 222. It's on page 275 of the PDF file.

MR. GEDDES: That is 6731. I'm sorry.

MEMBER BROWN: Which one are you talking about?

CHAIR APOSTOLAKIS: Two twenty-two. Event 222. This is the one?

MEMBER BROWN: I don't know.

MR. GEDDES: Yes, that sounds right.

CHAIR APOSTOLAKIS: Page 275 of the PDF. Is that the one?

MR. GEDDES: That's it, yes.

CHAIR APOSTOLAKIS: Good. Yes.

MR. GEDDES: Exactly. We call that at the

subsystem level because it affected one of the trip functions. The other trip functions were not affected by this defect in the reactor protection system.

The CPC in and of itself is a subsystem of the RPS.

CHAIR APOSTOLAKIS: Would it make sense to say -- you can go back to your slide.

MR. TOROK: Okay.

CHAIR APOSTOLAKIS: I mean you do a calculation here. You say, out of 27 common defect events, one could have resulted in a common-cause failure, and the ratio is 3.7 percent.

If I think in terms of common-cause failure models, would this be the beta factor? In other words, if there is a defect in one channel, the condition or probability of the same defect appearing in the other channels is .037.

MR. BLANCHARD: This is Dave Blanchard from AREI, and I participated in some of this classification.

No, I think when we are talking about the identical channels that have the same software and could be potentially subject to the same trigger, we are talking about a beta factor of one as opposed to .03. There might be a .03 chance of getting the

trigger perhaps.

CHAIR APOSTOLAKIS: No, because they say, out of 27 common defect events, one could have resulted -- okay? So the defect was common to both channels.

MR. BLANCHARD: Yes.

MEMBER BROWN: In the rod slip event.

MR. BLANCHARD: Yes.

CHAIR APOSTOLAKIS: It was already there.

MR. BLANCHARD: Yes.

CHAIR APOSTOLAKIS: So the trigger made the difference?

MR. TOROK: The problem here is that it goes back to the fact that the system misunderstood the phenomenon here. The system thought that a significant rod slip was going to take more than half a second, and that was built into the design throughout on the requirements. So, when there was a rod slip that happened in less than half a second, the system didn't recognize it. Right?

So it had nothing to do with the fact that the system was implemented in digital technology. The basic understanding of the phenomenon didn't recognize that that could happen, that the rod slip could be that short. That is really what drove the event.

MR. GEDDES: Of course, if it were an analog system, it would have, pretending for a moment that you could build a core protection calculator out of analog components, the problem was in the calculations that resulted in the subpoint itself, the parameter, which is an engineering process independent of the system design. Okay?

MR. TOROK: Now when we --

CHAIR APOSTOLAKIS: I am trying to understand what this 3.7 percent means. I agree with Dave that this is not a condition of probability of finding the defect, but then what is it?

If I am dong a PRA someplace and I am desperate for numbers, what does this number mean to me?

MR. TOROK: We were not trying to imply that --

CHAIR APOSTOLAKIS: I know you were not.

MR. TOROK: -- this was a number for PRA.

CHAIR APOSTOLAKIS: But now, moving one step ahead, Ray -- (laughter) -- the moment you put that number up there, you know, I get excited.

(Laughter.)

MR. TOROK: These PRA guys like numbers.

CHAIR APOSTOLAKIS: That is necessary but

not sufficient.

(Laughter.)

MR. BLANCHARD: As a PRA analyst, I am not sure that I could use the 3 percent number. I think the way they used it in the context of the OE report was appropriate. They looked at the different bins of failures that occurred in the digital systems, and they classified them and came to a conclusion about how much software common-cause failures contribute as compared to all the other causes.

So this is sort of a relative ranking of the different kinds.

CHAIR APOSTOLAKIS: Okay. So if I'm doing a PRA and I have a number for all other causes, then I can increase that number by 3.7 percent and say I have now included software, too.

MR. TOROK: That's creative.

(Laughter.)

CHAIR APOSTOLAKIS: Well, let me ask you this: why can't I do that? I mean I have a number here. I believe in your evaluation. So, boy, somewhere there, either in this report or in another report, I think it was you, EPRI, that says the contribution for software should be -- what? -- one or two orders of magnitude lower than everything else. I

believe that is EPRI.

So here I have now a way of actually trying to meet that, and you are saying no.

MR. TOROK: Well, we are saying we haven't thought about that.

CHAIR APOSTOLAKIS: Oh, that's what you're saying.

MR. TOROK: Yes. We didn't look at it that way. All we were trying to do was just get a handle on what fraction of the common defect events that we found were affected by the software, were controlled by the software. That is all we tried to do.

CHAIR APOSTOLAKIS: Would that be a good thing to do then, maybe not today, but to think about what that number, how that number would be useful in a quantitative evaluation?

MR. TOROK: We need to think more about --

CHAIR APOSTOLAKIS: Do you remember that in your report somewhere, not this one, but in another one, of the cost/benefit report, or one of those, you say that the software contribution should be one or two orders of magnitude lower than everything, the contribution from all other causes?

And I say, well, gee, that sounds

reasonable, and they give me a number here. So I can start saying something about it.

Look, I accept it, that you are not looking at this from that perspective, but maybe that is something you want to do in the future.

MR. TOROK: We've got it on our list now. We thank you for that.

MEMBER STETKAR: The only thing I caution you about, I think Dave has a good point, is that if you head in that direction, you are presuming that every single challenge in the world has an equal likelihood of --

CHAIR APOSTOLAKIS: I agree. I'm just asking. I see a number.

MEMBER STETKAR: Can we tease something useful out of that number?

CHAIR APOSTOLAKIS: Yes.

MEMBER BLEY: Can I tease something useful, but not necessarily --

CHAIR APOSTOLAKIS: I have not seen numbers yet anywhere that would be helpful to a PRA person. So now I see one.

(Laughter.)

MEMBER BLEY: And you're leaping on it.

(Laughter.)

CHAIR APOSTOLAKIS: What am I going to do with it? That's really what I am saying.

(Laughter.)

MR. TOROK: We certainly will look at that harder. Right now, I think we are all afraid to give an answer one way or another.

CHAIR APOSTOLAKIS: Are you going home tonight?

MR. TOROK: Pardon me?

CHAIR APOSTOLAKIS: Are you going home tonight?

MR. TOROK: Tomorrow night.

CHAIR APOSTOLAKIS: So you can give us the answer tomorrow.

(Laughter.)

MR. TOROK: Dave and I are going to go out to dinner tonight and we're going to have a beer and we are going to decide --

CHAIR APOSTOLAKIS: I think it takes a beer.

(Laughter.)

MR. HECHT: The two cautions are, No. 1, how sure are you that this is a complete listing of the relevant events? I mean I know you didn't advertise this as complete. I know you have two

sources, but there might be many other events which are happening that aren't recorded.

MR. GEDDES: In the nuclear power industry or --

MR. HECHT: In the nuclear power industry. There might be annoying things that are happening that nobody is bothering writing down.

MR. GEDDES: True.

MR. HECHT: Okay. And No. 2, maybe the way you should do that, if you want to do that, is by comparing non-software-based systems to software-based systems, and then doing the comparison that way.

MR. GEDDES: That's an interesting question. We kicked that around, and we thought we could go find how many times has an analog system resulted in a CCF, for example --

MR. HECHT: Right.

MR. GEDDES: -- and compare.

MR. HECHT: Right.

MR. GEDDES: But now we are talking about a much bigger effort that, frankly, just wasn't on the scope of the project at the time.

MR. HECHT: It wasn't in the scope of the project --

MEMBER BLEY: Of course, there is a big

report by Idaho on that very topic.

MR. HECHT: Well, what I am just suggesting is that, because of the very limited -- which is fair; you can define things. So long as you define your terms, it's fine. But because of the very limited way in which you define software failures, it might be directly reflecting what George was -- one of what I call, one of the Holy Grails, not that I think we all --

MR. TOROK: I think that is a good point: how do we get a handle on the non-software-based systems here? When we looked at it briefly, it became obvious that the number of events was going to overwhelm us relative to our resources for the project.

But, going back to maybe this Idaho study and other things, maybe there is a way we can get a handle on it. So that is another thing on our list. So thank you.

MEMBER BLEY: I am just kicked off. Bruce, you're from a utility.

MR. GEDDES: I have been, not currently.

MEMBER BLEY: Not currently? Okay.

I don't know how one would get a chance to even chase this, and there must be problem reports

that get sent back to the vendors of things that are driving people nuts along the lines of what you said.

MR. GEDDES: Yes. Yes, I think that's right.

MEMBER BLEY: So that might have a wealth of useful information in them.

MR. GEDDES: Yes.

MEMBER BLEY: But I don't know that those can be accessed in a reasonable way.

MR. GEDDES: I think LERs and OE reports, of course, LERs have very --

MEMBER BLEY: LERs have really -- you don't have to report a lot of things.

MR. GEDDES: Exactly. OE reports, I wouldn't say they're voluntary, but there's a certain sense of reportability shared by all the INPO members. Of course, INPO assesses the effectiveness of the reporting mechanisms and the effectiveness of the root-cause analyses, right? It is all about preventing events and sharing knowledge.

So I think there's a lot of information out there, but you're right, underneath every LER or OE report, there's a big, fat, thick file of all the information --

MEMBER BLEY: And somebody who would

really like to tell you about it probably.

MR. GEDDES: Well, I have been on root-cause teams where I have developed those files. Some of these events I have been personally involved in determining the root cause.

You're absolutely correct, there's much more information underneath, and it helps. When I read these reports, I can see the context.

MEMBER BLEY: You've been there. Exactly.

MR. GEDDES: Then, of course, we collaborate on the meaning and the taxonomy and all those discussions.

But you're absolutely correct.

MEMBER BLEY: I guess I don't personally know the threshold for reporting into --

MR. GEDDES: Into INPO?

MEMBER BLEY: Yes.

MR. GEDDES: Well, it's --

MEMBER BLEY: I'm sure the learning organization itself is encouraging people to report more and more.

MR. GEDDES: Yes.

MEMBER BLEY: But we don't know what the --

MR. GEDDES: If there is an event, it will

get reported. If it is a discovery of an issue, loss of function, for example --

MEMBER BLEY: Which is your last column, not your fourth column?

MR. GEDDES: That may not be a reactor trip.

MEMBER BLEY: Certainly, most likely.

I'm sorry. Go ahead.

MEMBER STETKAR: No. I was just going to tell Ray to be careful about his paper. You're covering the microphone, for our recorder.

MR. TOROK: Oh, sorry.

MEMBER STETKAR: It makes a lot of noise in the headset.

MR. TOROK: Maybe I should just cover that up there, just in case.

CHAIR APOSTOLAKIS: Be careful, yes.

MR. TOROK: Okay.

CHAIR APOSTOLAKIS: Okay, let's move on.

MR. TOROK: Let's try the next slide now, yes.

MR. GEDDES: Okay. Now I think this is where things got interesting the last time we were here, and we ran out of time.

(Laughter.)

It seemed like everybody was really interested in this information.

MR. TOROK: We should have started with this slide, you know.

(Laughter.)

MR. GEDDES: We felt like it would be useful to take these four events, 1, 10, 13, and 221. They are software-related the way we have defined software, and we tabulate the root cause, the failure mechanism, the failure mode, and the system-level effect.

Now, remember, our CCF idea is at the system level. That is where we draw our distinction. We have found subsystem-level CCFs, but they are not as interesting as the system-level CCFs in this research. Okay? This is where we spent a lot of time and energy.

This first event -- oh, then we included some taxonomy introduced by the ACRS. You guys wrote a letter, April 29th of last year, saying, wouldn't it be nice if we could go after these failure modes? We feel like there's some discussion about whether each of those things is really a failure mechanism or a failure mode.

We are going to talk more about that in a

few slides. Okay?

But we call these failure mechanisms where, for example, the first one, ACRS would tag that -- we believe this event would meet what you guys thought would be a task incorrect response. The value, the incorrect substitute value for failed sensor, that is the one event that we have spoken of as an example. Okay?

The failure mode would be a single channel may not trip when required, not the whole system.

MEMBER BROWN: Why did you call it "substitute"?

MR. GEDDES: In this case --

MEMBER BROWN: It is not deliberately substituted. It is a fallout of a failure, a sensor failure.

MR. GEDDES: No, I think the logic for that event, I would have to go back and look, but the logic for that event, if there is a failed sensor, the system would have inserted a substitute value, maybe a range clamp at the top end or the bottom end. I don't remember which way it might have failed.

MEMBER BROWN: A default value.

MR. GEDDES: Correct, a default value, correct.

MEMBER BROWN: For that condition.

MR. GEDDES: And in the logic, they discovered that the default value was not the right value. The effect, the failure mode would be at channel level, that the single channel may not trip, and at the system level, we classified that as a no CCF.

MR. HECHT: Can I just make an observation? I would say that the task incorrect response is kind of a category, if you want to say it for the specific failure modes, but I wouldn't call the task incorrect response to be a failure mechanism. A failure mechanism was that -- there was a specification error that is a class, and the specific instance of this class was that somebody put in, typed in the wrong value.

MR. GEDDES: Well, we kicked that around. In fact, we make that very clear when we get to Thuy's discussion of failure mechanisms and failure modes. In this case, we felt like that's really the root cause, and the failure mechanism that could lead to the lost channel would be an incorrect value substituted for a failed sensor. That is just the way we did it.

MR. HECHT: Well, the failure mode is

really something that you can observe externally.

MR. GEDDES: Yes, a failed channel.

MR. HECHT: Okay, but the incorrect response would be a more general statement. Well, actually, I would call that a no response, but that would be, to my mind, the failure mode.

MR. GEDDES: You see the problem?

MR. HECHT: No, I don't see the problem. Failure mode is a behavior.

MR. GEDDES: The distinction between a failure mechanism and a failure mode is something that we are prepared to discuss later in this presentation.

MR. HECHT: Okay, fine.

MR. GEDDES: We debated amongst ourselves, what do these terms mean? How are they applied? We felt like there was some confusion there.

CHAIR APOSTOLAKIS: I like the statement, though, that the failure mode is behavior, something you see, right?

MR. TOROK: Right, and in this case, we said the channel didn't trip, but we don't see what's going on inside there.

MR. GEDDES: In fact, this failure mode didn't actually happen. It was a defect that was discovered and reported.

MR. HECHT: But the failure mode would have been --

MR. TOROK: That the channel doesn't trip.

MR. GEDDES: Single channel failing to trip.

MR. TOROK: That's why we did it then.

MR. HECHT: Which is either incorrect response or no response.

MR. TOROK: Okay, you could call it no response, yes. That's right, you could have said that. Okay?

CHAIR APOSTOLAKIS: Okay.

MR. GEDDES: The next event, No. 10, is incorrect logic in the self-test mode. We call this a tasking correct response. The root cause was a design error, meaning the software design itself had an error. Okay?

In this case, the self-testing features in the system actually blocked safety injection. As the self-test was performed in a particular channel at a time, during that self-test certain portions of that test blocked safety injection. This was a four-channel or four-train safety injection system, sequencer system. The self-tests among each train were scheduled independently.

Dave did some calculations and found that, roughly, 15 percent of the time safety injection was blocked entirely.

MEMBER BROWN: So they were running asynchronously.

MR. GEDDES: Yes.

MEMBER BROWN: And even in that asynchronous mode, because of this design-toting software design error, then you come up with, that's how you come up with the 15 percent?

MR. GEDDES: And the way the self-test was scheduled.

MEMBER BROWN: Yes, okay. I got that.

MR. TOROK: But 15 percent in this case was enough for us to say, yes, that is a potential common-cause failure.

MR. GEDDES: That's a big deal, yes.

MR. TOROK: Yes, that was a big deal. If it had been, you know, a millionth of a percent or something like that, we would --

MEMBER BROWN: Well, no, that's why I was trying to pull on the 15 percent, because most at least the stuff I am familiar with -- "most" is the wrong word. The self-test we used in the programs for which I was responsible occupied 5 milliseconds out of

a 50-millisecond timeframe. And if you looked at the time to complete a self-test of that channel, that function, through all of its things it is supposed to do, we would be talking minutes to do that, two minutes, five minutes, depending on the complexity of the functions which you are checking.

So you have to factor in where that little piece is amongst that 5-minute overall period as well as the fact that it is running -- because they are going to be running asynchronously, they are not all going to get there at the same time. It is more than just not getting there at the same time one out of four, but you've got this significant amount of time relative to the time for that little piece to be tested.

Did that get factored into your 15 percent?

MR. BLANCHARD: Yes. Actually, what happened is that there were like 15 tests that were performed over the course of 16 hours. This was we just cycled through these. Each of the channels was -- they weren't synchronized.

But when you began a test, the safety injection signal could get blocked, but it didn't clear until the next test started, which would be an

hour later.

MEMBER BROWN: Until the next test started in that channel or another channel?

MR. BLANCHARD: That channel.

MEMBER BROWN: Okay.

MR. BLANCHARD: Okay?

MR. GEDDES: We have a slide that describes this event in a little bit more detail.

MEMBER BROWN: Oh, okay. That is a crappy design.

(Laughter.)

MR. BLANCHARD: Yes.

MEMBER BROWN: How did that ever get through?

MR. BLANCHARD: We classified this as relatively significant. Okay?

CHAIR APOSTOLAKIS: A more civilized term.

(Laughter.)

MR. BLANCHARD: And we've even taken it to the point of highlighting it here, so we can talk about it.

MR. TOROK: Yes, that's why this one is highlighted here. This one was special.

MEMBER BROWN: Yes, yes. I can understand that, even with my limited brain power.

CHAIR APOSTOLAKIS: Okay.

MR. GEDDES: Okay. Event 13 is about a radiation monitoring system process or lockup. If there is a momentary power interrupt, power would come back, but the processor would remain locked up without a clear indication that it was locked up. Okay?

And this was an RMS processor that could isolate -- I think this also had some control functions in isolating an auxiliary system, maybe an HVAC system. I would have to go back and look.

But the root cause was there's a missing requirement to have a watchdog timer.

MEMBER BROWN: A hardware watchdog timer.

MR. GEDDES: Correct. Yes. Software watchdog timers are generally not a good idea.

There was a WRITE operation in the software that was also a defect. So we considered this also a software defect. Okay?

So the combination of a missing watchdog timer and defect in the WRITE operation resulted in a task no response. In other words, it was locked up.

In this case, the trigger would be a loss of power, which would be considered a single random failure, and therefore, no CCF.

MR. HECHT: I might call that a hang, by

the way.

MR. GEDDES: A hang? Okay.

MEMBER BROWN: As opposed to?

MR. HECHT: No response.

MR. GEDDES: Well, you must have the definitions of the terms then.

(Laughter.)

CHAIR APOSTOLAKIS: Would you like to have them? We didn't include them, I don't think. We just had the list. It was supposed to be a trigger.

(Laughter.)

MR. TOROK: That's okay because we can come back to these mechanisms later.

CHAIR APOSTOLAKIS: If later we get them, you know, sure. I thought we gave a reference, though. We gave a reference.

MR. TOROK: Well, this is what we thought, events.

CHAIR APOSTOLAKIS: I think we gave more than one, in fact, if you go to the ACRS data we gave. But, anyway, if you have it handy, I'm sure the members would be interested, too.

MR. GEDDES: As you pointed out, some of these events could fit one or more of the terms.

MR. HECHT: Yes, and they may have to be

adjusted in this context.

MR. GEDDES: Yes.

MR. TOROK: And sometimes we are not sure precisely what the mechanism is, right?

MR. HECHT: I want to say that this is the mode.

MR. TOROK: Okay. Then you get to the mode and the effects, right? We will come back to that.

CHAIR APOSTOLAKIS: Yes.

MR. GEDDES: Okay. The fourth event on this table is 221. It is another radiation monitoring system, and a momentary step change in the output. This did isolate, I think it was, a containment ventilation system or an aux building ventilation system.

There's a spurious actuation, and a counter in the system was not initialized at the right time. That was a design error. The root cause was that there was an error in the design of the software itself.

CHAIR APOSTOLAKIS: Explain why isn't event 222 here. Wasn't that the common-cause failure? I mean, am I missing something here?

MR. TOROK: Was that in the safety system?

CHAIR APOSTOLAKIS: Yes. Wasn't it? Yes.

MR. TOROK: I don't know. Do they want to go back to 222?

MR. AUSTIN: Two twenty-two is listed as a non-software common-cause failure.

MR. GEDDES: I forget. Which one was 222?

CHAIR APOSTOLAKIS: The one we discussed earlier, page 275 of the PDF.

MR. GEDDES: That was a parameter error, not a software error. That is why it is not in this table.

Just go back to the slide.

MR. TOROK: Okay, okay.

MEMBER BROWN: We are calling it a power loss, rod --

CHAIR APOSTOLAKIS: No, it's on page 275 of the PDF, yes.

MR. GEDDES: Right. Where the lightning hit the transformer and the rod slipped more than they thought. There was a calculation that said, how much should a rod slip in a given amount of time?

CHAIR APOSTOLAKIS: Yes.

MR. GEDDES: That became a parameter, not a software design -- it wasn't in the code itself. It was an external number that you key in as a parameter,

like a tuneable parameter.

CHAIR APOSTOLAKIS: But go to the previous slide then, 13.

MR. TOROK: Right. The software design was fine. It was -- oh, wait a second.

MR. GEDDES: It was the constant that was incorrect.

MR. TOROK: Yes.

CHAIR APOSTOLAKIS: If you go to slide 13, in your box there at the top, you say, "Out of 27 common defect events, one could have resulted in a common-cause failure."

MR. GEDDES: That is event No. 10.

MR. TOROK: That is why this one is highlighted.

MEMBER BROWN: That was No. 10?

MR. TOROK: That's No. 10. That is the one where, had there been the trigger, the software safety function wouldn't have happened.

CHAIR APOSTOLAKIS: Okay. Okay. That's fine. That's fine. All right.

MR. GEDDES: Next slide. This is event 10. We have talked about it. We've classified it as a system-level potential CCF.

Do we need to discuss this any further?

MR. TOROK: We pretty much already addressed everything here.

MR. GEDDES: The next slide, I will turn it over to Dave. He looked at the risk significance of the event itself.

MR. BLANCHARD: Yes. I think, for every one of the class 1E events, in addition to classifying the failures themselves, we went through and did a risk significant determination. A lot of them occurred before the risk significant determination process existed, but we were able to apply the existing risk significant determination process to a number of them.

This is the most significant one that was found, and it is the one that we call the potential common-cause failure. The issue, of course, the significant determination process has this stair-step diagram where, as you go to the right on the diagram, it is increasingly more significant in terms of risk, in terms of an operating event or a failure.

Down on the left side of the chart, you have different initiating events that are considered as a part of the significant determination process from highest frequency to lowest frequency. Then, across the top, you have different levels of defense

in-depth and diversity that are considered in the significant determination process.

The issue here with this event was that, for 10 to 15 percent of the time, we would not be able to generate a safety injection signal, should an initiating event cause that trigger.

The concern, of course, was with the large LOCA, where there wasn't a lot of time for the operator to provide a backup to the safety injection signal. This still got classified as green or the lowest category in terms of safety significance because the safety injection signal was in the 10 percent range. It was available 90 percent of the time, which is roughly the same probability of failure as you would have with recovery of a failed train in the significant determination process. So perhaps it is right on the border between being in the green to the white in terms of risk significant.

But, in the risk significant determination process, we don't just look at a single event. We look at the whole spectrum of events that might occur. As we go in increasingly higher frequencies in terms of the initiating event, it is the medium LOCA, up to the small LOCA, now the operator can play more and more of a role in backing up the safety injection

signal. So the amount of defense in-depth and diversity is going up, and the "X" there is moving to the left. We have more defense in-depth and diversity as we have events that have lower and lower frequency and more and more time available to the operator.

Where we end up in the white area, for this particular event, is for the steam generator tube rupture. It has a high enough frequency and sufficient -- well, it has basically two different ways the operator can deal with a steam generator tube rupture, and at the same time backing up the safety injection signal.

These two trains of diverse mitigating systems here keep the core damage frequency for this particular event fairly low. However, the frequency of the steam generator tube rupture is high enough that in the significance determination process it would have been in the white category for this event.

MEMBER STETKAR: Dave, at the risk of just excruciating detail, this was a sequencer that the fault blocked the SI signal.

MR. BLANCHARD: Yes.

MEMBER STETKAR: Does the same fault also block the loss-of-power sequence?

MR. BLANCHARD: Well, as it turns out --

MEMBER STETKAR: And if it did, would the safety significance be altered quite dramatically because of the different frequency of that trigger?

MR. BLANCHARD: Yes, I think this was an unusual event in that, even though it was a sequencer for the diesel generators, it affected the safety injection signal when you didn't have a loss of outside power.

MEMBER STETKAR: Okay. Only part of it?

MR. BLANCHARD: Only that part of it.

MEMBER STETKAR: Okay.

MR. BLANCHARD: Which meant that you couldn't reduce the significance of --

MEMBER STETKAR: I'm just curious. Sometimes, especially with significance, if you look at a very, very specific event and only that failure --

MR. BLANCHARD: Right.

MEMBER STETKAR: -- you know, you might come up with a different safety significance --

MR. BLANCHARD: Right.

MEMBER STETKAR: -- determination if the same type of failure happened in a completely analogous circuit.

MR. BLANCHARD: Right. It turns out that,

had you had your DBA, had you had the LOCA, and the loss of outside power, the system would have worked fine.

MEMBER STETKAR: And if you had had only a loss of outside power, it would have worked?

MR. BLANCHARD: It would have worked fine, yes.

MEMBER STETKAR: Okay. So it was strictly that one --

MR. BLANCHARD: It was --

MEMBER STETKAR: Okay, fine. Thanks.

MR. BLANCHARD: It was during the testing mode.

MEMBER STETKAR: Enough detail.

MR. BLANCHARD: All right.

MR. HECHT: Go ahead.

MEMBER BROWN: No, no. No, go ahead. I will follow up.

MR. HECHT: I am having trouble understanding the columns, and I'm just wondering if you can help with a couple of questions.

MR. BLANCHARD: Sure. The significance determination process, basically, looks at the number of trains or systems you have available.

MR. HECHT: What is the distinction

between a train and a system?

MR. BLANCHARD: A train might be a single-train system that would be subject to a single failure, such as -- I don't know. In a BWR, a HPCI system is just a single-train ECCS system. All right?

MR. HECHT: So I am saying greater than three trains or two redundant systems. So does that mean that the systems are doing different or have different functions?

MR. BLANCHARD: Yes, that's correct. The two redundant systems, let's take, let's say, a loss of feedwater. You have two redundant systems available in the first in the form of auxiliary feedwater, which is a multi-train system. And backing that up, you would have the ability to do cooling feed-and-bleed, which would be a safety injection and PORV set of systems. So that would be the two redundant systems in that column.

MR. HECHT: So, when you say something like one train plus recovery of failed train, what does that --

MR. BLANCHARD: Okay. In the case of the steam generator tube rupture, what we are talking about there is the one train would be the ability to equalize pressure between the steam generator and the

reactor and terminate the leak through the tube. That would basically be a single-train system.

A recovery of the failed train would be manual actuation of the HPCI system, given that the safety injection signal had been blocked by this failure.

MR. HECHT: I see. So this is actually a combination of the static design plus the states of the systems?

MR. BLANCHARD: Oh, yes. The significance determination process takes a look at the event, puts all the systems in the state that they were in at the time of the event, and then you look at the diversity and defense in-depth, given that plant condition.

MR. HECHT: I see.

MR. BLANCHARD: All right?

MR. HECHT: Thank you.

I'm sorry.

MEMBER BROWN: But you asked one of my questions. So that worked out okay.

MR. HECHT: Okay.

MEMBER BROWN: So this is the event where you had one hour if something gets blocked, one hour before it restarts or it is reinitialized or to trigger the reset of the test.

MR. BLANCHARD: Whatever reset it.

MEMBER BROWN: And I've got some number of channels or trains that I am dealing with, and I see nothing is ever in the red. In other words, I'm interpreting your table to show that there would never have been a circumstance where you did not have a response, safety injection performance in time to mitigate the downstream effects of a LOCA or a stuck-open relief valve.

MR. BLANCHARD: No, on that, what red means is that there is sufficiently little defense in-depth available, that this event becomes risk significant. I believe like the righthand, where you see the red in these columns is around 10 to the minus 6, isn't it, for the event?

I think the threshold between red and yellow is around 10 to the minus 6 per event.

MEMBER BROWN: Okay. So, since you never get to something, all of these are less than 10 to the minus fifth in the lefthand column?

MR. BLANCHARD: Yes.

MEMBER BROWN: Does that mean --

MR. BLANCHARD: Well, perhaps the threshold at green is at 10 to the minus 6. What you are seeing here is everything here is 10 to the minus

6 for each one of these initiating events with the possible exception of the steam generator tube rupture, which might be slightly higher than 10 to the minus 6 for a steam generator tube rupture with this particular condition.

MEMBER BLEY: I think we can go through this another time.

MEMBER BROWN: Yes.

MR. BLANCHARD: This is a whole-day discussion.

MEMBER BROWN: Just a quick break here. We were supposed to stop for lunch here.

MEMBER BLEY: I just noticed we're about a fourth of the way through the slides.

MEMBER BROWN: Yes, but I mean I was looking for a stopping point about four slides from now to get through. Do you want to do it now and just pick up the conclusions and the non-1E events after lunch?

We've got about six slides to get through, if you want to get through them all before we get into the failure modes and stuff like outside the conclusions.

Your druthers is happy with me. We've got to get through slide 22, and we're not making a lot of

slide-per-minute progress.

MR. SIEBER: At our rate, it will take another hour.

MEMBER BROWN: Pardon?

MR. SIEBER: At our rate, we would take another hour.

MEMBER BROWN: Yes. So my suggestion, my option is to go ahead and go to lunch now, come back, and start doing the rest of these.

MR. GEDDES: May I suggest that we have one more slide on 1E, and then we can break for lunch --

MEMBER BROWN: That's fine.

MR. GEDDES: -- and come back?

MEMBER BROWN: Okay. That's good.

MR. GEDDES: Okay, next slide.

Mr. Hecht I think had already picked up on this slide.

We show that, for example, incorrect parameter values is more frequent human error than others. We don't propose to throw out these events, that we have programs and processes and root-cause analysis and corrective actions that go after these kinds of things. But the point is some of these things can be equally applicable to analog systems.

We have policy and rules and guidelines that help us address those mechanism.

Okay. So that is what the OE tells us. That is where some of these events fall out. That's the only point here.

MEMBER BROWN: Okay.

MR. SIEBER: Let me ask a quick question, and it would probably just take a sentence or two to answer.

One of the issues that I have had to do with this EMI issue, and it was on the DC power supply to digital devices, what I learned was the characteristics for EMI change with time.

MR. GEDDES: Yes.

MR. SIEBER: Particularly if you have contacters and other things on there that, when they are new, put out pretty clean changes in the power supply, but as they age and operate, all of a sudden, the arcs become longer, and the EMI effects become larger.

If you go to a plant and test it, and it passed with flying colors, how do you take into effect the age-related changes in EMI effects on input buses? And how does a licensee --

MEMBER BROWN: That is a one-sentence

answer?

(Laughter.)

MR. SIEBER: Okay. Is that the answer?

MR. GEDDES: Well, we found that if a component was susceptible to a no EMI condition, noise or a fast transient, we have seen several events related to electrical fast trains, as to a relayed kickback, for example, on an input signal.

MR. SIEBER: Right.

MR. GEDDES: That is a common defect. That means that equipment is somehow, by design or age-related degradation mechanisms, or some other means, susceptible to EMI. So we classify that as a common defect.

However, they are usually manifested in the form of a single failure. Okay?

MR. SIEBER: Yes. It can be.

MEMBER BROWN: As long as you have separate power supplies on that channel --

MR. GEDDES: Yes, and adequate separation on our cables.

MR. SIEBER: And providing that the aging is occurring at different rates.

MR. TOROK: And this is 1E. So you've got the separation of your power, and so on.

MR. SIEBER: Right.

MEMBER BROWN: Does that satisfy you then? Well, not satisfy you, but are we finished?

MR. SIEBER: I am finished.

(Laughter.)

MEMBER BROWN: Okay. All right. Okay, we will go ahead and close out or adjourn the meeting for one hour. We will be back here at -- adjourn, I'm sorry, at 12:00; suspend, I'm sorry. Suspend.

George will not be here. So you will have to put up with me at one o'clock.

(Whereupon, the foregoing matter went off the record for lunch at 11:51 a.m. and resumed at 1:02 p.m.)

A-F-T-E-R-N-O-O-N S-E-S-S-I-O-N

1:02 p.m.

CHAIR APOSTOLAKIS: Okay, we are back in session.

Which slide are you on, Ray? My God, you're going back?

(Laughter.)

MR. TOROK: This is the last one we got through. We're done with this one.

CHAIR APOSTOLAKIS: Okay.

MEMBER BROWN: This might engender a comment about we've got 62 minus 18 slides to go.

CHAIR APOSTOLAKIS: This is the full presentation from you? All the stuff that is on the agenda?

MEMBER BROWN: Yes.

CHAIR APOSTOLAKIS: I thought this was just --

MEMBER BROWN: Thirty-four more slides to go through.

(Laughter.)

CHAIR APOSTOLAKIS: Well, I will leave it up to Mr. Torok to manage his time. You're so slow, Ray.

(Laughter.)

MR. TOROK: Thank you. Well, I appreciate the opportunity to try to manage my time with this group.

(Laughter.)

As you know, it is a challenge.

CHAIR APOSTOLAKIS: I'll tell you. I think you should go straight to the main messages. You know now where the Subcommittee is coming from. They get excited when they see data, the evaluation of data, how did you do this, and all that, and possible conclusions, of course.

MR. TOROK: Right.

CHAIR APOSTOLAKIS: So, although I suspect most of your slides are of that nature from now on --

MR. TOROK: We tried to do it that way.

CHAIR APOSTOLAKIS: Yes.

MR. TOROK: There's some discussion of the failure modes and effects, and whatnot, in the middle, in between the OE stuff and the PRA stuff, just because there are linkages there we wanted to establish.

CHAIR APOSTOLAKIS: It's up to you, Ray --

MR. TOROK: So let's just go on it.

CHAIR APOSTOLAKIS: -- and your colleagues.

MR. TOROK: Very good. So Bruce is going to resume, and we're now on non-1E system common defect events.

MR. GEDDES: Right. I only have a few slides, and then we are done with OE -- theoretically.

(Laughter.)

CHAIR APOSTOLAKIS: There is a lot of bitterness here.

(Laughter.)

MR. GEDDES: I said that with good cheer.

(Laughter.)

CHAIR APOSTOLAKIS: This is a happy family, right?

Please, Bruce.

MR. GEDDES: Okay. Three hundred and twenty-two events, this time 273 are non-1E events. Out of those, 77 we found a report that said there was a common defect, and 20 of those due to software, 57 non-software. The same taxonomy, the same structure, except non-1E systems tend to lose their independence at one point or another. Okay? Non-safety systems, we have a slide that we think explains the differences.

The key point on this slide is that we found seven CCFs, meaning both redundancies were

affected, triggered simultaneously, due to software, 9.1 percent out of 77. Okay? That's the data.

Next slide.

CHAIR APOSTOLAKIS: So you guys are going to think about the meaning of those 9.1 percent?

MR. TOROK: Yes, that is on our list.

CHAIR APOSTOLAKIS: Very good.

MR. TOROK: Thank you.

MR. GEDDES: Now let's look at software failure mechanisms. This is actually the last slide we showed you the last time we were here, and it prompted the most excitement and the invitation to come back.

So you have seen this slide before, except we replaced the word "failure modes" with "failure mechanisms" because our colleague from EDF --

CHAIR APOSTOLAKIS: Keep going.

MR. GEDDES: Okay. Well, we broke down these software. We recast the title as failure mechanisms, and you can see the breakdown.

This table, this Pareto chart, is no different from the one we showed you last time.

Eight of the 20 were related to application logic errors, buffer overflow. That could be an operating system or platform issue or it could

be an application coding issue. You could see it different ways.

CHAIR APOSTOLAKIS: Again, I don't know. I am stalling here until Myron shows up.

(Laughter.)

You crossed out "modes" and wrote "mechanisms"? That is, when you guys were doing the ATHEANA stuff, you spent a lot of time thinking about what is a mechanism in fact and what --

MEMBER BLEY: For us, the mechanisms were the things that went on inside the head.

CHAIR APOSTOLAKIS: The mode is the manifestation of an error?

MEMBER BLEY: It would be, but we didn't actually use the term "mode".

CHAIR APOSTOLAKIS: You used what, something else?

MEMBER BLEY: Human failure event.

CHAIR APOSTOLAKIS: Human failure event?

MEMBER BLEY: And unsafe acts. Failure mechanism --

CHAIR APOSTOLAKIS: But Myron said earlier about the failure mode is observable.

MEMBER BLEY: Observable, yes, the way he has categorized them.

MR. TOROK: Actually, that's sort of our next topic, right after we do away with the OE.

MR. GEDDES: We can explain what we mean very clearly and succinctly when Thuy's portion of the presentation comes up.

MR. TOROK: But, basically, we gave the presentation last year. There was a letter that came out, I think with some of Myron's input there, that talked about understanding of failure modes, and there was a list of them, and so on.

CHAIR APOSTOLAKIS: Yes.

MR. TOROK: So we went back and looked at what we had after that and said, wow, we called those failure modes, but we should have called them mechanisms, and we had better go explain why.

CHAIR APOSTOLAKIS: So that was as a result of our letter?

MR. TOROK: That's what we are working up to here.

MEMBER BLEY: I think it is fair to say, from the Subcommittee's point of view, and I might get knocked down in a hurry, I don't care so much what you call it; I want to understand what went wrong. That's what we're after.

CHAIR APOSTOLAKIS: Yes, but I mean --

MR. TOROK: But there's a reason that it is important to understand mechanisms and modes and effects, and we have a bunch of material on that.

MEMBER BLEY: Okay.

MR. TOROK: Do you want us to go ahead with that?

CHAIR APOSTOLAKIS: Yes, go ahead.

MEMBER BLEY: What if we align this picture, the one on the 1E failure mechanisms.

MR. TOROK: Here, let me drive for a second. Where's PageUp/PageDown?

MR. GEDDES: This slide, because there are four, we don't give you Pareto chart; we give you the events. Because there is only four, we can put them in a table and examine them in some detail.

MEMBER BLEY: Right.

MR. GEDDES: Down here, now we say, what were the others, the other 23? Now we make a Pareto chart, so we can see how they rank next to each other in terms of frequency. Okay?

MEMBER BROWN: So the Pareto chart is the bar graph?

MR. GEDDES: Yes. This bar graph is about non-1E software failures. There's too many to put in a table on one slide. So we make a bar graph out of

it. Okay?

MEMBER BROWN: Got it.

CHAIR APOSTOLAKIS: Actually, maybe you're right, these are mechanisms, and they are consistent with human error and knowledge.

MEMBER BLEY: Yes.

CHAIR APOSTOLAKIS: It's not what's going on, but may lead to some fault, yes. So it's okay. It's okay.

MR. GEDDES: Okay. We can go on?

CHAIR APOSTOLAKIS: Yes.

MR. GEDDES: We spoke earlier about the inherent design differences between 1E and non-1E systems.

MEMBER BROWN: Before you go on, you can answer my question, I guess. Is there a distinction -- I'm trying to get a distinction between modes and mechanisms.

CHAIR APOSTOLAKIS: Do you agree with slide 19?

MEMBER BROWN: While he is looking at that, a mechanism is something that starts, but the mode is the mode of failure that it takes after the mechanism? Am I on the right terminology?

CHAIR APOSTOLAKIS: Yes.

MEMBER BROWN: Okay, thank you.

CHAIR APOSTOLAKIS: If I didn't sleep at all at night and then I screw up in the morning due to something, the mechanism is the lack of sleep.

MEMBER BROWN: Okay. I just wanted to make sure -- I'm trying to understand your terminology. I want to get it right.

CHAIR APOSTOLAKIS: I think the actual failure mode is that I did something to my machine here.

MR. HECHT: But there is a additional concept which we have to introduce.

I apologize. It just took much longer to make a reservation change than I thought it would.

But the additional concept is really the level of indenture, if you will, the level at which the analysis is being done. So, for example, a failure mode, if we are just talking about the computer, might be different than the failure mode if we are talking about what's happening at the system level. So you have to define both.

MR. GEDDES: That's right. That's right. Right. That's how we view it. One man's failure mode is another man's failure mechanism, depending on where you are on the hierarchy of the system.

MR. TOROK: And actually, we are going to come to that in a few minutes. All we are doing here was we took a slide that we had shown you before and said, wow, those things that we called "modes" the last time we should have called "mechanisms", and let's talk about why we should have called them.

MR. HECHT: I have to tell you one more thing. That is that I am not sure that mechanism, which I would call a cause, if we were talking in the DoD world, but we have this first -- we have the mode and then we have first-level effects, second-level effect, and then an effect.

So I just wanted to clarify that it is not so simple as saying one person's cause is another person's mechanism -- I mean one person's failure mode is another person's failure mechanism. I think you have to define the level at which you do it, define the failure modes that are appropriate for that particular level, which I think is basically the computer, if we get down to it.

MR. GEDDES: Yes.

MR. HECHT: And then speak about what's happening at the next-level effect, which might be the train or the division, and then the end effect, which might be -- or the next-level effect, which might be

the system level, and then the end effect, which might be --

CHAIR APOSTOLAKIS: But all these effects could be failure modes of the subject subsystem.

MR. HECHT: But then define it, then define the failure modes which are appropriate to the level at which you're doing it.

CHAIR APOSTOLAKIS: Yes, but the use of the word "mode" is appropriate.

MR. HECHT: Yes. Yes, it is.

CHAIR APOSTOLAKIS: It's just that you make a sequence of effects.

MR. TOROK: Yes. It depends on what you're doing. If you are designing systems and components, the mechanisms and modes are of interest. If you are designing the big picture of the plant, then the modes and the effects are of more interest.

CHAIR APOSTOLAKIS: Okay. So we all agree then that the change in terminology of this slide is appropriate?

MR. GEDDES: There's a NUREG that --

CHAIR APOSTOLAKIS: Oh, then it should be right.

(Laughter.)

We agree, right, Myron?

MR. HECHT: I'm not sure if faulty dead-band function -- what does that mean? Was that an incorrect response? Was that --

MEMBER BROWN: It's a dead band that is either too long or too short.

MR. HECHT: So that might be --

MEMBER BROWN: That's a mechanism.

MR. HECHT: No, that would be too late, early or late response, right?

MR. GEDDES: We took the ACRS failure modes --

MR. HECHT: I see.

MR. GEDDES: -- as described in your letter and applied them; we only applied them to those four events. We didn't take it to these additional 20. We could, but we didn't.

MR. HECHT: Okay. All right. Well, I am not sure that these are mechanisms.

MEMBER BROWN: Can we work on that later?

MR. HECHT: Okay.

MEMBER BROWN: So we can get through this stuff?

CHAIR APOSTOLAKIS: Good. The point was made. Let's move on.

MEMBER BROWN: Thank you.

MR. TOROK: Okay. Moving on --

CHAIR APOSTOLAKIS: So you are on 20, yes.

MR. GEDDES: Twenty, yes. We have presented this slide before, and it is worth just recapping.

There are inherent differences in the design and the design criteria for 1E versus non-1E systems, and this goes to how perhaps triggers can influence how a CCF can come about.

Mr. Brown, I think you mentioned that independence is one of the strongest features in a 1E system.

MEMBER BROWN: One of the biggest belts of armor.

MR. GEDDES: Correct. Thank you. We agree.

In a non-1E system, we tend to see more master-slave architectures, which means only one controller can be operating the final component at a time. So, at some point, there are shared components and single-point vulnerabilities, and by definition, those are common defects that can be triggered into a CCF. Okay?

So we draw this distinction in the data because the underlying criteria we believe help us use

the OE to better inform how we develop solutions for particular applications.

So, in terms of actual potential CCFs where there are common defects, we saw 6 out of 27 for 1E systems and 38 out of 77 for non-1E. The fact that the non-1E systems were more than twice should not be a surprise because of the inherent nature of those non-1E systems. Okay?

CHAIR APOSTOLAKIS: So, again, separating the trigger from the defect, when you say, "actual or potential CCFs", you mean without the trigger or with the trigger?

MR. GEDDES: An actual CCF is one that is triggered.

CHAIR APOSTOLAKIS: It has happened?

MR. TOROK: It includes the effect of triggering, what you're talking about --

MEMBER BLEY: And the demand, from what you said.

MR. TOROK: Yes. The problem -- well, first --

CHAIR APOSTOLAKIS: So the 6 out of 27 includes -- I mean the thing was, in fact, demanded and failed or potentially --

MR. TOROK: It means there were common

defects and the potential for common triggers. Regardless of whether they actually happened, common triggers could have happened, and therefore, it was at least a potential CCF.

Now we lumped actual and potential together. We said they are of equal import for our purposes.

CHAIR APOSTOLAKIS: But maybe in these six, in one case the trigger actually happened; in others it didn't. So the trigger effect is there.

MR. TOROK: We showed you that before, right? For the 1E systems, there were no events where there was --

MEMBER BLEY: Concurrent.

MR. TOROK: -- concurrent triggers actually happened. There was one where they --

CHAIR APOSTOLAKIS: But that's why I am asking the question. If I look at this number now, does it include the occurrence of triggers or the potential occurrence of them?

MR. TOROK: Both.

CHAIR APOSTOLAKIS: Okay, great. Great.

MR. TOROK: Both, and the big difference here is that non-1E systems tend to share parts, and that makes you much more vulnerable to common-cause --

CHAIR APOSTOLAKIS: But, I mean, don't you think, though, that for 1E systems 22 percent is a pretty high number? I mean, should we be scared here?

MR. HECHT: It all depends how often it happens, right?

CHAIR APOSTOLAKIS: From the PRA perspective, if I see a number like .22, I know you're going to think about it, but now you have included the trigger there, the potential for occurrence of the trigger.

What am I going to do with that? Is that my common-cause failure rate?

MR. TOROK: We have lumped two things together here. One is the existence of the common defect, and the other is the existence of --

CHAIR APOSTOLAKIS: The trigger.

MR. TOROK: -- triggers.

CHAIR APOSTOLAKIS: Well, thank you very much. So, if I have a system, that is some measure of the probability of failure. Maybe it is .22. Maybe it is something else.

MEMBER BLEY: But this is of cases in which you had common defects.

CHAIR APOSTOLAKIS: It is conditional, like the beta factor.

MEMBER BLEY: Yes, yes.

CHAIR APOSTOLAKIS: Is this my beta factor?

MEMBER BLEY: No, no.

CHAIR APOSTOLAKIS: Why not?

MR. BLANCHARD: Excuse me. This is Dave Blanchard.

This is not your beta factor. We have several categories of common-cause failure here. Some of them are software-related; some of them aren't.

If we knew the number of operating hours or the number of demands, the denominator, from these numbers, we could figure out the probability of occurrence of common-cause failures. Then we could partition those common-cause failures into software-related common-cause failures and non-software common-cause failures.

The 22 percent or the 27 percent, whatever it is, that is not the beta factor. The beta factor is the number of common-cause failures over some denominator, which right now is undefined. We don't know how many demands there have been, nor how many successes there have been. We don't know how many operating hours there have been.

MEMBER BLEY: We don't even know the

population.

MR. BLANCHARD: We don't know the population, right. But if we could get that information from these numbers here, we could distribute the common-cause failure probability, whatever that was, between software and non-software common-cause events.

MR. NGUYEN: If I may add my grain of salt? Initially, most people think that if there are common errors, there is systematically common-cause failure. This data shows that in one incident the fact that you have these common errors in different channels does not mean necessarily that there would be common-cause error.

In the one case of the failure, there might be an actual or potential common-cause failure, but in four cases out of five the fact that there are common errors does not mean that there will be common-cause failure.

MEMBER BLEY: Can I take you back to something else?

MR. NGUYEN: Yes.

MEMBER BLEY: These two pictures, the classification scheme we have on slide 19 -- you don't have to jump to that -- for the non-1E software

failures, if we tried to apply that classification scheme of what you called mechanisms to the events on page 14, I can see one for sure that fits into one of those and one that might. Have you done that? I mean we have different names for similar kinds of things.

MR. GEDDES: We haven't done that exercise. We could, sure.

MEMBER BLEY: Okay. So we don't have a common set of names, bins into which we're partitioning these things?

MR. GEDDES: Well, we do in a certain sense. I read the reports. If a report said there's an application logic error, there's one. If I found one report, though, I got two --

MEMBER BLEY: And one of these was about here.

MR. GEDDES: Right.

MEMBER BLEY: That is the only one I saw that I could clearly align between the two.

MR. GEDDES: Well, there is a logic error in event No. 10 that is --

MEMBER BLEY: That is the one I'm --

MR. GEDDES: Right. That one is similar -- I would consider that similar to some of the logic errors we found in non-1E systems --

MEMBER BLEY: Okay.

MR. GEDDES: -- in the nature of the defect. Okay? Meaning that is at the application level, not at the operating system, not buried pieces of modules, but the actual program that makes the system perform the usual function.

MEMBER BLEY: Right. Okay. It doesn't quite look like any of the other three quite aligned with the ones that occurred in the non-1E. Is that true or am I missing the boat there? It isn't completely clear.

For comparing the two sets of things and making conclusions, you are doing that with numbers, but it would be nice to also be able to do it with the kinds of failures that occur.

MEMBER BROWN: You mean the specifics don't line up with any of the other four?

MEMBER BLEY: Yes, but I am not sure of that. If I had a common set of bins into which I would group failures when I find them, then I could better compare things between one kind of system and another.

MEMBER BROWN: Well, the application logic error --

MEMBER BLEY: Yes, we talked about that

one.

MEMBER BROWN: Yes, it's similar. I mean you could argue that it's --

MEMBER BLEY: That one fits.

MEMBER BROWN: But the rest of them --

MEMBER BLEY: I'm not sure.

MEMBER BROWN: Yes, it's not quite obvious at all.

MR. GEDDES: I think where we strike a difference is, for example, event 10 is a potential CCF, and in a non-1E system it might actually be a CCF.

MEMBER BLEY: Sure.

MR. GEDDES: Because it is constantly under demand.

MEMBER BLEY: But the mechanism that happened, that's what I was trying to get at.

MR. GEDDES: Right. Right.

MEMBER BLEY: What we have called mechanisms on the other picture. We are not using the same categories of things, looking at the two kinds of systems.

MR. TOROK: And I think that is because the categories came right out of the OE reports and LER reports. We are using those words as opposed to

inventing our own set of bins and then trying to put them in it.

MR. GEDDES: Right. Exactly.

MEMBER BLEY: I can understand that, but from a point of view of reassembling all this into something useful, it seems to me you need to bridge that gap.

MR. GEDDES: Yes, I see what you are saying.

MR. GEDDES: That is a good observation.

MR. TOROK: Yes.

MR. HECHT: Okay. If I were to just look at that 22 percent in that case, the reason why, for example, a logic defect didn't result in a common-cause failure or didn't have the potential for a common-cause failure is because the sensors were different or because a channel was in a maintenance state or something like that.

But had the sensor data been the same two multiple channels, then it would have been the same result on multiple channels.

MR. TOROK: Yes. In that case, yes.

MEMBER BLEY: The trigger was the failure in the sensor, right, not that --

MR. GEDDES: If you failed two sensors in

the application logic in the same way, and the application logic has a defect that results in an incorrect response to a failed sensor, then both channels would be affected.

MR. HECHT: So that, I guess, emphasizes your point that there are a lot of things that have to go wrong in order for a software defect to cause a disaster?

MR. GEDDES: Well, the recipe for a CCF is a common defect and concurrent triggers.

MR. HECHT: Yes.

MR. GEDDES: We want to examine and attack both of those problems.

MR. TOROK: Is it okay to go on?

MR. HECHT: Yes. Just one more question. I'm sorry.

You didn't consider the voter in any of these situations. Because, in actuality, of course, there is a voter.

MR. GEDDES: Well, we only considered the events on the systems that were reported.

MR. HECHT: Yes.

MR. GEDDES: If we didn't see a report that called into question the voters or how the voters should have or would have behaved, then it won't

appear in the dataset.

MR. HECHT: Yes. Okay, but I think it is important to observe that, ultimately, even in these situations where you think you have dependence, you don't.

MR. TOROK: Yes, that's a good point.

MEMBER BROWN: Well, it depends on where the voter is executed also and how. If it is executed in the software, in the program, of if it is executed outside the program, we are looking at some type of voting hardware. Whether it be solid-state switches or a combination of logic units, or what have you, that's one. But when you are doing that voting inside the program loop and it is a subroutine, that introduces its own complication or potential to be affected.

MR. HECHT: But, still, ultimately, you have one control logic --

MEMBER BROWN: Yes, yes.

MR. HECHT: That's in that funnel or, you know --

MEMBER BROWN: Oh, yes, the output of the voter is -- but you've already had the problem by then.

MR. TOROK: And you are right, in

evaluating the potential for common-cause failure, that is certainly a valid consideration.

I'm thinking we're okay now on this. The only thing I wanted to come back to, just very briefly on this, is the point of this slide really for us was, when you try to combine the 1E data with the non-1E data, that is problematic, and there are some good reasons why.

MEMBER STETKAR: Let me ask you about that because this is a lot of -- you've identified five attributes, and your assessment of those attributes qualitatively reinforces your observations.

What I would ask is, you've actually spent a lot of time examining the real events that have happened and thinking about whether we call them failure mechanisms or failure modes or failure causes, or whatever bins we throw these things into, among these five attributes, which is the most important attribute that makes non-1E systems so much worse than 1E systems?

MR. GEDDES: Independence.

MEMBER STETKAR: Independence?

MR. GEDDES: Or lack thereof at some point in the system.

MEMBER STETKAR: So it's the first one,

the redundancy issue?

MR. GEDDES: Yes.

MEMBER STETKAR: Okay.

MEMBER BROWN: Well, it's independently redundant. You can have redundancy, but you're not necessarily totally --

MR. GEDDES: If you share a power supply or a sensor, which a lot of non-1E systems will do --

MEMBER STETKAR: Okay, but that has nothing to do with the digital system and it has nothing to do with software. It has everything to do with system design. It could be analog.

MR. GEDDES: Yes.

MEMBER STETKAR: It could be two valves headed off the same piping system.

MR. GEDDES: That is true, but --

MEMBER STETKAR: It has nothing to do with what we are looking at.

MR. GEDDES: No, it does, because in the reports, if the word "digital" or some variation of that keyword search, resulted in a hit --

MEMBER STETKAR: That's a keyword search. I'm thinking about, what did you think about it? Well, but that's okay. That's a keyword search, and many people would throw those out as saying, well,

obviously, if I have a common AC power supply for two motor-operated valves, if that AC power supply -- that has nothing to do with the valve design. It has nothing to do with the size of the torque switches on the motors. It has nothing to do with anything. It is not relevant to the issue that I am examining.

MR. GEDDES: We actually started down this path, and we didn't include it in the final report, but we do believe digital systems give you additional fault tolerance, if you choose to implement fault tolerance, like monitoring power supply outputs, and if the first power supply fails, additional systems can do a better job of telling you if you choose to take advantage of that kind of feature available in the technology. There are lessons learned still in those digital systems.

MEMBER STETKAR: Stepping back now from the independence of redundance, which of the other four then are the next largest contributor, or can you do that? And it's okay if you say, no, that you haven't really thought about it; that's fine.

I was just curious whether -- what I am trying to think of is that you claim that 1E systems are always very, very simple. They are always very independent. They always have no interaction with

anything else.

Well, suppose I am looking now at an 1E application that, indeed, I don't know what you mean by simple, but does not seem all that simple to me. Is that something that triggers my sensitivity to looking at specific issues, that the non-1E, for example, experience is more relevant in that particular area?

MR. GEDDES: I would say, first of all, there are cases where 1E systems can be complicated. I mean core protection calculators are more complex than simple functions. These are general observations.

MEMBER STETKAR: There are probably even foreign applications of some of the integrated protection and control systems that you haven't looked at that perhaps Thuy is more familiar with that are even more complex.

MEMBER BLEY: It's kind of what I was trying to ask in the other way. I mean, given that you've got a common defect, we have kind of three times as many of the bad actors in the non-1E systems, but I was trying to look at the failure mechanisms between the two and saying, is this because one kind of failure mechanism occurs a lot more over in these

non-1E systems, or is it something else?

Why aren't I seeing the same kind of things that I see in the 1E systems causing failures in spades over in the non-1E systems? I'm a little confused by not seeing what I would expect that way.

I don't know if that makes sense to you or not.

MEMBER STETKAR: I was just trying to step back to this: there are some of these attributes --

MEMBER BLEY: Without combining the data, and we can see lots of reasons why you wouldn't want to do that. Can we learn by combining inferences from the two things?

MR. GEDDES: I think so. You know, if there is an application defect, then quality assurance methods and V&V become very strong tools to defend against that, whether it is safety or non-safety.

We do much more formal V&V and formal reporting for 1E systems, but non-1E systems are coming along. Plants are learning from these events, and giving V&V, for example, or formal software quality assurance methods much more respect. Equipment reliability and plant operations are big drivers.

MEMBER BLEY: These things cost you money.

MR. GEDDES: Yes. I could tell you, over the last few years, reporting and sensitivity and equipment reliability have really increased. Twenty years ago, the quality of the event reports didn't give us nearly as much information as they do today.

It is no accident that our capacity factors have improved by 10 percent over the last 10 to 15 years. It is a lot of this type of equipment reliability, and digital helps us, if we implement it correctly.

MR. HECHT: I have on the previous slide one last question.

MR. TOROK: Did you want to go back?

MR. HECHT: I can't go back, can I? Okay.

You have formal SQA methods and you say "always" and "varies". What proportion of plant digital control systems are purchased as commercial products?

MR. GEDDES: Control systems?

MR. HECHT: Yes.

MR. GEDDES: They are always purchased as commercial-grade items.

MEMBER BLEY: You mean kind of off-the-shelf items?

MR. HECHT: Yes, as off-the-shelf items.

In other words, the same --

MEMBER BLEY: You mean a feedwater control system is an off-the-shelf item?

MR. GEDDES: No, I wouldn't put it that way. I think the modules that make up a feedwater control system, the controllers, the I/O modules, the buses, the --

MEMBER BROWN: The individual assemblies.

MR. GEDDES: Those can be catalog items or they can be manufactured to a spec, but, generally, they come from a commercial source. Okay?

MEMBER BROWN: Right. So this is an area where the plant operator doesn't really have control.

MR. HECHT: Right.

MR. GEDDES: No, the plant operator can specify and insert himself in this process and have as much control as he would like. He doesn't have to install this equipment. If he's not satisfied that some level of quality has been achieved, he won't install. Nobody installs -- maybe I'm being a little too, I don't know. I don't think anybody would install a system with known defects.

MEMBER BROWN: They would never install a system with no defects?

MR. GEDDES: No, "known".

MEMBER BROWN: Oh, "known"? I'm sorry.

(Laughter.)

MR. HECHT: That might be true, but I don't know whether Allen Bradley or Foxboro individual control modules, I don't know how they have been designed.

MR. GEDDES: Well, what we mean by varies and improving is that the OE, plants look at the OE from across the industry, not just our own corrective action system. This same report is being now used in the form of case studies, for example, that teach engineers why SQA methods are important, and they are now inserting themselves more today than they were five or ten years ago because of the interest in equipment reliability.

Some vendors, you know, quite frankly, say, "You guys are really being a pain in the neck. Nobody else does it like this." And a lot of my colleagues would say, "So? That's the way I want it." And if you don't want to sell me your services, the integration, application engineering services, I'll go somewhere else or I just won't do the project."

It's not worth the event. These events are very painful, and the root-cause process that comes out of that causes -- you know, people's

behaviors change when the pain of change is less than the pain that they are in, right? It is easier to change than to fight.

So that's what we do in our root-cause process. We are trying to get behavior changes across to engineers. They can, in turn, influence vendors. That's what's making that improve.

This OE, in some ways it is embarrassing on the non-safety systems. We shouldn't have that many events, but we are learning from them. That's why we are improving. That's all this is really trying to say.

MR. HECHT: Okay. Thank you.

MR. GEDDES: Conclusions, insights, and inferences. You've heard us say that software has been no more problematic than other contributors. You've heard us say it is difficult to combine 1E and non-1E experience, and why we believe that.

You've heard us say there's no events for diverse platforms. In other words, the specific instance of platform diversity would have been effective in protecting against CCF.

We have found several events where a loss of one function in a protection system did not result in the loss of other functions that would come into

play for a wide range of transients. For example, containment pressure versus pressurizer load, and protecting against CCF.

Okay, next slide.

Recommendations. There is something we are doing right, and we want to examine that in more detail and reinforce it going forward. That doesn't mean we are perfect. We have more to learn. We have a ways to go, but so far the trend is that other forms of CCF causes are more dominant, and we are already attacking those as well. So, whatever we are doing on software, we need to keep doing it and get better at it.

And everybody has mentioned we should get additional OE from other countries, nuclear countries and industries, to see if our results are consistent or if there's additional lessons learned that we can deploy in our fleet.

MEMBER BLEY: Before you leave this, I really liked what I have seen so far and what I saw the last time you guys were here. I think there are some areas where you can get more information out of what you've already done. I think there are places that somebody could mine to start maybe thinking about how you would model some of this.

I am curious as to what programs you actually have for doing more of these things. You were talking about, and I am wondering if the MOU has reached a point where you guys are talking about joint products, or is that going to come up in the next couple of days? Is somebody going to talk about that?

MR. SANTOS: Dan Santos, Office of Research.

The answer is, yes, we actually have a specific project called Operating Experience, and a component of that will be collaboration with EPRI. We are not that far along yet. EPRI is pretty recent. But we envision getting there throughout this year and early next year.

We will talk about it --

MEMBER BLEY: We would sure be interested in hearing about the plan for doing that.

MR. SANTOS: I plan to give you the details tomorrow and answer follow-on questions --

MEMBER BLEY: Great.

MR. SANTOS: -- on that specific.

MEMBER BLEY: That's great.

MR. HECHT: Given the point that you have there about the difficulty of combining 1E and non-1E experience, what would you say about efforts to look

at software failure experience from other industries?

CHAIR APOSTOLAKIS: You shouldn't even try.

MR. HECHT: Is that your --

CHAIR APOSTOLAKIS: They would be --

MR. HECHT: Well, we would certainly welcome the opportunity to look at that for the purposes of what we have done. We had access to a lot of information from the U.S. nuclear power industry. That's what we went with, but --

CHAIR APOSTOLAKIS: I think looking at other countries in their nuclear industry would make perfect sense. But, as Myron says, I mean if 1E and non-1E are difficult to combine -- now, if I go to railroads, I don't know.

MEMBER STETKAR: Now wait a minute. I know nothing about the aircraft industry, but if the aircraft industry employs simple, redundant digital controllers in their aircraft because they have decided that that's an appropriate thing to do, why wouldn't the experience from simple, redundant digital --

CHAIR APOSTOLAKIS: I think in that industry you have mostly control systems, not single like ours.

MEMBER STETKAR: But I don't know that.

CHAIR APOSTOLAKIS: Do the nuclear first.

MEMBER BLEY: Okay, but the underlying idea that, because you don't want to combine the data from non-1E and 1E means you can't use the two together to draw more useful information, see how to build models, and that, to me, might well extend to other industries. It might not. But once we have a good framework for looking at the 1E and non-1E together and understanding how they are related and not related, we might move forward --

CHAIR APOSTOLAKIS: Every single time the nuclear guys say they are going to look at other industries to learn and this, they learn nothing. That's what I am saying. We are a unique industry, being regulated to the point of pain, and you are going to go now somewhere else to learn? Good luck.

I would like you to go, though, to the Korean experience, Taiwanese, Japanese, French, Swedish, the nuclear, do that first. If to satisfy my colleagues you want to do the other stuff, fine. I am not going to object to it. I'm just saying I don't have high hopes we are going to get anything out of it.

MEMBER BROWN: We did -- I'm sorry -- you

know, we did look at -- from the diversity standpoint, we --

CHAIR APOSTOLAKIS: Or practice, what do they do?

MEMBER BROWN: They did, the research did look at a wide variety of industries. They are different, and they pointed that out in the study.

CHAIR APOSTOLAKIS: Yes. That's the Oak Ridge report.

MEMBER BROWN: But most of those systems were --

MEMBER BLEY: I guess, even before we go fetch foreign reactor experience, understanding how to categorize and use the information we have already collected seems to me an important first step. But go ahead.

CHAIR APOSTOLAKIS: Yes, sure.

Yes, sir?

MR. NGUYEN: I agree, more or less, with you.

CHAIR APOSTOLAKIS: Will you guys give us anything?

MR. NGUYEN: Oh, yes, of course.

(Laughter.)

On the other industries, there is one

industry for which it might be interesting to get information. It's the process industry. Because the platforms that we use for control systems in nuclear power plants are, in fact, the platforms that are used in other industries.

MEMBER BLEY: That's where they were pioneered.

MR. NGUYEN: Yes. So, in fact, these vendors, these platform vendors, in fact, have collected their own operating experience.

CHAIR APOSTOLAKIS: Guys, fine. Go ahead and do it. I am not saying don't do it. I don't control your resources. But I know what is going to happen.

(Laughter.)

Nuclear experience in other countries, though, it is really a great thing. It really is. Isn't it OECD, CSNI, ABCDEFG group that looks at INC, and they are doing something like the common-cause failure guys used to do, and they are still doing, in fact, collecting experience? Is there such a group?

MR. NGUYEN: There is. One of the big problems is the fact that the reports are in national languages.

MR. GEDDES: You just confirmed what he

just said.

CHAIR APOSTOLAKIS: Thank you.

MR. NGUYEN: Yes. So it's why the --

(Simultaneous speakers.)

MR. NGUYEN: -- needs to be done, I would say, as projects, not as --

CHAIR APOSTOLAKIS: Come on, Thuy. The guys who did the hardware common-cause failure exercise had the same problem, right? I don't think the Swedes write in English only for that, although the Swedes are pretty good; they do. Let's say other countries.

I think looking at the nuclear experience, maybe that can be the good conduit because it is an international organization. So, you know, as long as you don't come up with any conclusions, because then they're international, which means they mean nothing.

This is a really frank discussion.

(Laughter.)

MR. TOROK: We appreciate your comments. We have looked at some of this information from other industries, like aviation and so on, and I think my personal reaction is there are lessons to be learned there perhaps, even if we are just looking at what they do. They use the same techniques we do. They

just use a different mix of them.

If we can look at what they are doing, understand why they are doing it, then we come back to us and say, okay, we see why they're doing what they're doing, but that's not really right for us. But it helps us understand what is right for us. That is worth something by itself.

CHAIR APOSTOLAKIS: If you look at the transactions and reliability and all that stuff, do you know how many papers are out there on human reliability -- I mean software reliability models. Do you know many are useful?

MR. TOROK: Well, I could take a guess.

CHAIR APOSTOLAKIS: Minus 2 percent.

(Laughter.)

MR. TOROK: Okay.

CHAIR APOSTOLAKIS: I'm telling you.

MR. TOROK: Okay. Shall we move on? Good idea. That was a hint to move on, wasn't it?

CHAIR APOSTOLAKIS: I would like Charlie to see one of those papers and tell me how useful they are.

MEMBER BROWN: I stopped all my transactions subscriptions years ago, for the exact same reason.

CHAIR APOSTOLAKIS: For that same reason.

MEMBER BROWN: Well, I wasn't smart enough to understand them.

MR. TOROK: Now you guys are really going to like this next topic.

CHAIR APOSTOLAKIS: Okay, next topic.

MR. TOROK: This is where we try to take the next step. We told you we learned certain things about mechanisms and modes, and so on, from the OE. We want to talk more about that for a couple of reasons.

One has to do with this quote from ACRS. "Digital INC may introduce new failure modes that are not well-understood." And there was a list of items in that letter, and we are going to come back and talk about those some more.

The other thing we want to reference here is this fault tree handbook, NUREG-0492. Some of you may be familiar with that.

It turns out it is a really good reference on this topic. It explains relationships between mechanisms, modes, and effects.

CHAIR APOSTOLAKIS: Very good.

MR. TOROK: We think that is very applicable in what we are doing. Okay? So we will

get more into that.

So, in terms of subtopics, we want to tell you a little about what goes on right now in terms of FMEAs for real digital systems going into plants. FMEA is done by vendors, and so on, and Bruce is going to explain that to us.

Then we are going to talk a little more about what goes on inside the box in terms of realistic digital system behaviors -- this is Thuy's game -- and how that relates to the context of the nuclear plant, which, of course, leads us to, okay, so what's the "so what?" for PRA? Where did we get the numbers? So that is where we are going now. Okay?

CHAIR APOSTOLAKIS: Good.

MR. TOROK: So digital FMEA practice, Bruce, please take it.

MR. GEDDES: Oh, okay. I've seen a wide range of FMEAs and digital systems, ranging from thousands of pages to just a handful of pages, considering single failures. The IEEE 352, and I think there is a MILSPEC or a MIL standard, asks us to postulate single failures, determine the method of detection, if there is one, and then the effect on the system. Okay?

This is where mechanisms and modes

sometimes get confused. I have seen FMEA discuss both. I have seen FMEA discuss mechanisms only, and others discuss failure modes.

But, ultimately, we want to understand the impact at the system level. They are deterministic. The design engineers out there do not attempt to assign probabilities in their FMEAs. They are used more as a design tool to make sure we understand the mechanisms that can lead to failure modes, and try to design them out before we install the system or, as a minimum, make sure we have a clear method of detection via indications or alarms or a combination of both.

We have seen some good practice where, for example, one utility takes their FMEA as an artifact of their design process, and they develop troubleshooting tools for maintenance after the system is deployed. So, if they see a certain effect, they can work backwards and determine where the failure mode or failure mechanism occurred to improve troubleshooting.

MEMBER BLEY: Has that worked well? I have never seen anybody do that.

MR. GEDDES: That utility reported to me that that has worked well, and we have advised other utilities to adopt that practice.

MEMBER BLEY: And are they doing that massive kind of FMEAs?

MR. GEDDES: Well, these are non-safety systems. For example, the accounting FMEA is thousands of pages, where there's 18 cabinets and there's an appendix for each cabinet that exhaustively treats every component and other mechanism.

MEMBER BLEY: It would be hard to use that to generate troubleshooting.

MR. GEDDES: Exactly. But I've seen, for example, a digital feedwater system FMEA that went into 100-plus pages, and that could be useful for an engineer to help a maintenance guy generate a troubleshooting procedure or even a pull-tree troubleshooting pullout, for example, the back of a procedure.

MEMBER BLEY: Is that getting sent around the industry anywhere?

MR. GEDDES: Yes, yes. I have seen some lessons learned papers being distributed at some of the conferences.

MEMBER BLEY: It would be nice to find some of those.

MR. HECHT: It's a practice that is used in the defense industry.

MEMBER BLEY: The stuff I have seen over there was of such massive sort, I couldn't find its usefulness.

MR. HECHT: Well, what you do is, for example, you have an end effect, which is what the person sees. Then you just kind of sort the failure modes by end effects, and then you go to intermediate effects. In other words, it starts looking and seeing, given that end effect, which intermediate effect gets it. That reduces your number of original causes.

By the time you get to the third one, you can sometimes identify what it is, assuming that the engineers who have done the FMEAs have done their job, and, of course, assuming that there's only one thing that went wrong, which is generally not the case of unreliable systems. There are multiple things that go wrong.

MEMBER BLEY: It would be nice to see something really useful coming out.

MR. GEDDES: I can give you a paper that I wrote last year for --

MEMBER BLEY: I would be delighted if you could pass that on through Christine.

MR. GEDDES: Okay.

We have seen software functions credited for fault detection tolerance. That is the strength, it can be the strength of a digital system. That is how a digital system can give us more reliability than analog systems.

We think failure modes are well-understood at the system or even the component level. We are still learning about failure mechanisms. Okay? We put that footnote in here. We think the first time we put this together that saying that failure mechanisms are well-understood might be an overstatement. We are still learning. It is important. When Thuy explains what we mean by the difference between mechanisms and modes, you will see how and why that is.

For example, the taxonomy alone can be confusing. What's the definition of task no response versus a task incorrect response? What mechanisms lead to that kind of a result?

MEMBER BLEY: When you say they are well-understood, I would expect you to have had a really solid answer then when I asked you about comparing the failure mechanisms in the 1E and the non-1E systems, that, yes, we really understand this and we can tell you exactly what's going on.

So I am a little skeptical, and I don't

see that in the report, I don't think, of how that's so well-understood and available yet. If more is coming, that's great.

MR. GEDDES: Yes, more is coming in the presentation.

CHAIR APOSTOLAKIS: But, again, going to event 222, the delay in stopping --

MR. GEDDES: Okay.

CHAIR APOSTOLAKIS: -- was that a surprise or do you think you understand? We could have said, yes, this is one of the possibilities, or is that a delay in your terminology?

MR. GEDDES: In terms of a failure mode, it would be a delayed response.

CHAIR APOSTOLAKIS: A delayed response?

MR. HECHT: But in terms of --

CHAIR APOSTOLAKIS: That doesn't mean we understand it just because we call it that? I don't know.

MR. HECHT: Well, it depends on what you are using it for. The reason why you want to do the failure mode is so that you can come with detections and things. If you are trying to engage in a process of fault avoidance, then you might be looking at the root causes or the mechanisms.

MR. TOROK: Exactly.

CHAIR APOSTOLAKIS: Failure mechanisms and modes are well-understood. So that means if I were doing an evaluation, an analysis, let's say not BLA, but the first part, the event investigation I would have said there may be a delay here in scramming because I understand it well. Would I have said that?

MR. TOROK: No.

CHAIR APOSTOLAKIS: Is that something you would expect people to do?

MR. GEDDES: No.

CHAIR APOSTOLAKIS: So how well do we understand it? I mean, after the fact, we say, oh, yes, sure, that makes sense; this is what happened. The question is, a priori, when you are doing an analysis, do you understand them well enough to start listing possible failure modes?

MR. TOROK: For the late response, what you do is you have timers.

CHAIR APOSTOLAKIS: So then it comes naturally, you're saying? You will worry about it?

MR. HECHT: Yes, because your FMEA will say, for the reactor trip response, reactor trip function.

MR. TOROK: A lot of mechanisms are well

enough understood so that the designers are incorporating features to deal with them right now, and they have been for a long time. That is some of what Thuy will talk about.

CHAIR APOSTOLAKIS: But I don't think you can claim the same thing for feedback control systems. So these are different. I mean we are talking here about simple systems that are shutting down something. They start something else. They open yet another thing.

If I go to a complex system like the Arianne rocket which is automatically controlled, I am not sure you can make that statement, which is fine. You don't have to make universal statements, but let's not forget --

MR. HECHT: No, no.

CHAIR APOSTOLAKIS: I don't think so.

MR. HECHT: The Arianne V failure involved the two inertial reference systems shutting down. In other words, no response. That led to a loss of stability. The failure mark was quite clear.

CHAIR APOSTOLAKIS: But mine are after the fact.

MR. HECHT: But why did they shut down?

CHAIR APOSTOLAKIS: Look at it and say,

yes, sure.

MR. HECHT: No, the point is that -

(Simultaneous speakers.)

MR. NGUYEN: It seems to me that when we say failure modes, there's two of them, we might understand different things. What we can do is, for a reasonably simple system with reasonably simple functions, we can identify the possible failures. That's right.

CHAIR APOSTOLAKIS: Now that is not a universal statement if I include feedback and control systems, which we don't have to worry about at this stage anyway.

MR. TOROK: The plant does. Maybe the NRC doesn't, but --

MR. GEDDES: What we are finding in the OE reports, for example, is that there are interesting failure mechanisms that occur in control systems with dynamic memory allocation, for example. And you find out that the defensive measures that the integrator put in place for controlling data in and out of memory were not very robust, and a piece of information ends up in the wrong place, and now a PID control block is acting on erroneous data and making feed pumps.

CHAIR APOSTOLAKIS: So are you gentlemen

saying that this red statement there is correct, no matter what the system is?

MR. TOROK: Well, wait a second now. What you said is right in that the more complex you make a system, the more difficult it is to anticipate all the strange behaviors it might have.

CHAIR APOSTOLAKIS: Yes.

MR. TOROK: That's true. However, we would say that, for a lot of real live systems, for most real live systems, certainly real live systems going into safety applications where the functionality is simpler than that, then for the most part both the mechanisms and the modes are quite well-understood.

CHAIR APOSTOLAKIS: And that's what I'm saying, too.

MR. TOROK: They function, I would say, very well.

The mechanisms we will talk about a little bit, too. In fact, as I said, the designers right now, and for decades, the designers have been incorporating features to deal with specific mechanisms. They have been aware of --

MEMBER BLEY: And they are still doing that because we are still learning.

MR. TOROK: That's right.

MEMBER BLEY: That is why it seems to me we are on a learning curve. We are not up here at the top yet.

MR. GEDDES: We are, and we are talking about digital FMEA practice, what's actually happening in the field. Okay?

You pick up an FMEA for Oconee. It's exhaustive and it's proven. It's tested. It's repeatable. It's demonstrable in the factory acceptance test environment or in the integration environment.

If I pick up a five-year-old or a ten-year-old FMEA on a feedwater control system or a feed pump speed control system, I might see it is kind of cryptic, that there's been an event, that the FMEA didn't contemplate a failure mechanism inside the event that helped contribute to the event that was not postulated or understood. So we are improving. That is where we are.

MEMBER BLEY: Let me push you just a little further because most FMEAs I've studied, and I haven't studied FMEAs on INC systems, only look at independent failures. They don't look at interactions. There must be interactions here that are really important for us that we are beginning to

learn more about, too.

CHAIR APOSTOLAKIS: Right. Absolutely.

MR. SIEBER: But once they buy it, they think I'm done.

(Laughter.)

MEMBER BLEY: I think Jack hit it.

MR. TOROK: It is true that as an industry we are still learning how to do this better. In fact, there's an EPRI project that it looks like we are going to do next year in regard to this, because our members have basically said, "Look, we have put in these systems. We've had trouble with them that shows us that our FMEAs, which we actually did, were maybe not as good as they should have been."

On top of that, we get this 1,000-page FMEA. In real life, it is awfully hard to take advantage of that. Can't we focus on the stuff we really care about and do a better FMEA?

Well, wait until next year, and maybe we can come back and explain where we are.

CHAIR APOSTOLAKIS: First of all, I think the work you are doing is great. It really sheds light where we thought it was darkness.

(Laughter.)

I will compare it with something that I

think did the same thing many years ago, when people were saying errors in human behavior, misdiagnosis, and it was a whole misdiagnosis. My God, misdiagnosis. How did we do it? How did we do it?

Then a guy had a simple idea. He developed a little table and he said, well, what is the actual event that can be misdiagnosed as what? And that was a major step forward. It turned out it was only one or two things, you know, the small LOCA and the steam generator tube rupture.

And you look at it and you say, "My God, I was scared that things would be misdiagnosed and all hell would break loose, when in fact it's not that bad."

I think that is what you are doing here. This is a great step forward. It really is.

And if we seem to argue every now and then, it's our nature. We cannot help it.

(Laughter.)

MR. SIEBER: We still think of the red statement, though, as a goal.

CHAIR APOSTOLAKIS: Not the criterion.

MEMBER BLEY: Or a fact.

CHAIR APOSTOLAKIS: But, still, though, I think for simple command systems, you may be right;

maybe we are almost there. Yes, yes, that is what I am saying.

MR. SIEBER: Unfortunately --

CHAIR APOSTOLAKIS: Speaking of foreign experience, there is an interesting incident at the Bruce Reactor in Canada, which I would like you to evaluate. That was a control --

MR. TOROK: When did this happen?

CHAIR APOSTOLAKIS: Bruce?

MR. TOROK: Yes. When?

CHAIR APOSTOLAKIS: Like Bruce.

MR. TOROK: Yes, yes, we got it.

CHAIR APOSTOLAKIS: In Canada.

MR. TOROK: Recently?

CHAIR APOSTOLAKIS: Oh, no, it's been years.

MR. TOROK: Okay.

CHAIR APOSTOLAKIS: It was not a simple system. It was not a simple system.

MR. TOROK: Okay.

CHAIR APOSTOLAKIS: I don't remember the details now.

MR. TOROK: Okay.

CHAIR APOSTOLAKIS: If I try to remember them, I'll screw up.

Are you familiar with it? Somebody is. Yes, Mike is.

MR. GEDDES: I think I have heard about that.

MR. TOROK: Okay. Now let's wrap up this topic, okay?

In terms of FMEA experience --

MR. GEDDES: I think we have touched on all those things.

MR. TOROK: We have? Okay.

CHAIR APOSTOLAKIS: Okay. Let's keep rolling.

MR. TOROK: Okay. In that case, we are going to move along and talk about modes and effects. So Thuy --

MR. NGUYEN: We have already talked a little bit about that.

CHAIR APOSTOLAKIS: So look at their conclusion. "Digital system often has the same set of possible failure modes." Yes.

Now why do you put the word "often"?

MR. NGUYEN: Well, because --

CHAIR APOSTOLAKIS: Well, you know, this is a rare-event business.

(Laughter.)

MR. NGUYEN: Yes.

MEMBER BROWN: Digital systems have all the failure modes that analog systems have, and we add software into it, and we add additional functionality into it in terms of things to do. Okay? And we incorporate the potential for interactions from channel to channel, which adds additional complexity in terms of the failure mode.

So digital systems bring a lot more things that can go wrong, depending on how you decide to design or employ it.

MR. GEDDES: Or go right. Fault detection --

MEMBER BROWN: Which kind of right/left are you talking about here?

(Laughter.)

I'm already far enough right, according to most people.

MR. GEDDES: Oh, I'm with you.

(Laughter.)

CHAIR APOSTOLAKIS: We are talking about different systems.

MEMBER BROWN: But, as long as the fault detection is done in a manner in which you don't compromise, guess what?

MR. NGUYEN: The normal practice, when you specify a system, whether digital or not, is to, in fact, say what are the failure modes that come into this definition you can accept, and what are the failure modes you have to avoid as much as possible. So I would say the notion of failure modes is, I would say, something that, I would say, at least in practice is well-identified.

CHAIR APOSTOLAKIS: These three sub-bullets?

MR. NGUYEN: Yes.

CHAIR APOSTOLAKIS: Failure to actuate, late --

MR. NGUYEN: Yes. In the case of a simple --

CHAIR APOSTOLAKIS: Is there such a thing as premature or that's spurious actuations, right?

MR. NGUYEN: Yes, that's right.

MR. TOROK: Now you notice also for the purposes of the definition, what Myron said earlier was the failure mode can be thought of as the behavior viewed from outside the system, right? That is the same thing we are saying here.

MR. HECHT: And also, for your simple on/off failure protection functions, you are viewing

that not from the processor level, but you are viewing it from a system level.

MR. GEDDES: Right.

MR. HECHT: And for that reason, that statement that you haven't read is true. But if we were to look at it from, once again, the thing that software does is it implies a processor, and all digital systems have one or more processors in them. Then the balance of that control system --

MR. GEDDES: Failure modes at the processor level would be --

MR. HECHT: Would be quite different.

MR. GEDDES: -- would be mechanisms that lead to these failure modes at the system level.

MR. NGUYEN: So maybe we can go to the next slide.

MR. TOROK: Let's go on. Let's go on.

MR. NGUYEN: So, in a definition, the failure mechanism is an event or a chain of events that occur during operation and that leads to a failure.

So a mechanism is not necessarily a very simple thing. It could be a chain of events, starting at a very low level of a very small component and sneaking its way to affect the whole digital system.

MR. HECHT: I see.

MR. NGUYEN: Okay? And in fact, the different mechanisms could lead, or very often lead, to the same failure mode.

MR. HECHT: No, the same effect.

MR. NGUYEN: Again, to me, according to the previous slide, the failure mode is a behavior of the digital system as viewed from the outside.

MR. HECHT: Yes.

MR. NGUYEN: Okay? So, for example, if I have a stray radiation that modifies a memory cell and that leads to a spurious activation --

MR. HECHT: Yes.

MR. NGUYEN: -- it has the same, this failure mechanism leads to a failure mode which is spurious activation, and spurious activation could be closed by a very completely different failure mechanism.

MR. HECHT: Well, I would say that in your example there spurious actuation is when you have actuators which are causing physical phenomena to happen. But if we were to look at the processor level, what's really happening is you are getting an incorrect result.

MR. NGUYEN: That's right. That's right.

MR. HECHT: Okay, but do you understand the difference is that the way you have defined mechanism is that there's a chain of events that is happening. One of the events is that the computer output an incorrect result that led to transmission of a signal through a communication system to an actuator or maybe to a second computer that caused an actuator to be actuated.

But the point is that the failure mode is not the same as the effect and it is not the same as the mechanism.

MR. NGUYEN: Oh, yes, I agree with you. The effect is something else.

MR. HECHT: Okay.

MR. NGUYEN: Okay. So, in fact, the reason why we separate modes and mechanisms is because failure mechanisms are very technology-dependent and very dependent on the design. The objective of the designer is to avoid as reasonably as possible the failure mechanisms that could lead to, I would way, the failure modes that you want to avoid.

MR. HECHT: Oh, okay. Once again, you are saying that a mechanism causes a mode?

MR. NGUYEN: Yes, but --

MEMBER BROWN: And a mode has an effect.

MR. HECHT: No, I would -- oh, okay.

MEMBER BROWN: I'm a little bit lost on the --

MR. HECHT: On the distinction.

MEMBER BROWN: -- on the distinction because I can see a failure mechanism causing a mode of failure, which caused an effect of an actuation going on, which now causes flow to stop or causes rods to drop or causes, you know, the loss of flow in a loop. So that is the effect I look at.

MR. HECHT: Right.

MEMBER BROWN: The mode is what's generated by the mechanism that causes the control -- so we may plow that up and down. I just think we are getting wrapped around the axle on --

MR. HECHT: I think if you draw a picture and define your terms, that would help.

MR. GEDDES: This is what we mean, if you will just bear with us just for a second.

MEMBER BROWN: We just had a failure mechanism, which resulted --

MR. GEDDES: This was intentional. This is from the fault tree handbook, NUREG-492. Some of you I believe were involved in the preparation of this NUREG back in the day.

This is an example of a fuel flow system. If you look at corrosion of actuator stem, bottom left corner, corrosion of actuator stem, and go across, at the actuator --

CHAIR APOSTOLAKIS: Can you move a little bit, so we can see better?

MEMBER BROWN: Which one are we looking at?

MR. GEDDES: We are looking at this table, the bottom row, where it says, "Corrosion of Actuator Stem". Now go across. That's a mechanism that leads to a failure mode of the actuator. Okay? Because the actuator stem has binding. That binding is a failure mode. The effect is the valve is unable to open.

So you play the same game moving up in the hierarchy. As the binding of the actuator stem happens at the valve, that is a mechanism that leads to the failure mode of the valve to not open. The effect is no flow.

And you keep repeating this game. So we are talking about system versus component, controller versus -- it's in this context that we are trying to make these points.

MR. HECHT: What this table shows is another instance of the point that I had made earlier.

MR. GEDDES: You did.

MR. HECHT: You have to define the level of indenture, and maybe an important part of this, and maybe the point of confusion, is tell us what --

MR. GEDDES: What we mean?

MR. HECHT: Tell us which one of the rows of that table you are talking about for digital INC systems, and I think that would be universal across digital INC systems. I think it would be the processor interface.

MEMBER BROWN: I don't agree with that. I mean if you have an -- well, I don't know. I'm just trying to take what you said, and I can have a mechanism of a -- I'll take your previous one, the sensor failure, the trigger, whatever it is. That is a mechanism, isn't it? I mean that happened. It converted something to a mode somewhere.

The failure mode is you got incorrect data into something, and the processor couldn't handle it. So that would create a mode of failure in terms of it wasn't going to generate the proper sample or algorithm processing, or whatever. The effect was it told something to not operate when it should have.

MR. GEDDES: In a single channel.

MEMBER BROWN: In a single channel.

MR. GEDDES: You've got it. That is it.

MEMBER BROWN: That is the way I -- so I was trying to connect the dots. So the processor is not part of the chain, but it occupies a different place in the overall chain.

MR. TOROK: I think it depends on what level of abstraction you're operating on. If you are trying to design a better digital gadget, then you want to understand the mechanisms so that you can design features that can help you avoid them.

If you are trying to model the system in PRA, you don't care about that part. You care about the effects at the plant level and the failure modes perhaps. So it depends on the level of abstraction that you are operating.

MR. HECHT: But isn't the focus of this work on digital INC systems?

MR. TOROK: Yes.

MR. HECHT: If this were an analog system, you wouldn't care. I mean the analog system has to respond to a bad sensor as much as a digital system does.

MEMBER BROWN: I would look at that the same way.

MR. NGUYEN: The only point, the reason

why we introduced the failure mechanisms is because we want to have a sufficiently good understanding of the mechanisms so that we can prevent them from occurring.

MR. GEDDES: Put together a better design.

MR. NGUYEN: Yes.

MR. TOROK: Let's skip ahead a few slides here.

Actually, in this presentation we are kind of operating at multiple levels of abstraction really because we talked about OE. That's what is going on in the plant, that this is the level.

We are going to talk about what's inside the box here in terms of mechanisms and modes. We are going to come back to the system plant level that PRA -- we operate at different levels of abstraction here.

Anyway, what I was thinking was we talked about a factsheet, and I think we can just skip this one. What I wanted to get to was this list from that letter from April of last year, I guess. And it characterizes these things as modes, as failure modes.

We started looking at this saying, wait a second. For our purposes for most of what we are doing here, are these modes or are these mechanisms?

Can we go to the next slide?

MR. NGUYEN: So, in this slide, I have, I would say, summarized the principles, the design principles, of many digital reactor protection functions, where after an initialization phase the computer enters an infinite loop, usually of a fixed duration, typically, 50 milliseconds.

At each cycle, the software reads the inputs in sequence. So it has, I would say, so many input modes and so many communication ports. It will read them one after the other and will put what it has read in predefined places.

After that, the computer will execute the application code, which will read the inputs, do whatever it needs to do, and compute the values that will lead upwards to the higher modes.

So, after the execution of the application, the software will retrieve the results, the application results, and will run them on the output boards one after the other.

So this is, I would say, repeated at each cycle. After that, usually the software has not exhausted the 50-millisecond cycle time, and it will perform some auto-tests, until the limit of 50 milliseconds is reached, and then it will start again.

When it writes the values on the output

boards, automatically the hardware, that will reset a watchdog timer. If the watchdog timer doesn't see a new write on the output boards within 50 milliseconds, or usually slightly longer, it will decide that the software has entered some unknown state. It will shut down the computer and generate an output value, which is usually a signal which could be either a spurious shutdown or an analog load, whatever.

And this is repeated every 50 milliseconds. So there is no notion of tasks or there is only one task. There is no, I would say, sharing of the processor and of the memory by multiple tasks running in apparent concurrency.

MEMBER BROWN: The reset, you talked about writing results to the output, which I agree with your picture, except you have the reset of the auto-tester done within the 50 milliseconds.

MR. NGUYEN: Yes.

MEMBER BROWN: Yet, you have the reset to the watchdog occurring after the board outputs the write, instead of completing its entire thing. But, yet, you said the hardware reset occurs after longer than the 50-millisecond mechanism fixed cycle time.

So that I didn't understand your diagram because there was an inconsistency relative to when

the hardware timer resets.

MR. NGUYEN: Yes, the watchdog is reset by the fact that the output values are written on the output boards. The watchdog expects that 50 milliseconds late at least --

MEMBER BROWN: Okay. You start the 50 milliseconds starting on performing the auto-test?

MR. NGUYEN: That's right.

MEMBER BROWN: Okay. All right. That's fine.

MR. HECHT: Well, I would say that that may or may not be true, and I would say that in most cases I would find that extremely difficult to implement.

MEMBER BROWN: What, this?

MR. HECHT: Yes, but let me finish. Very difficult to implement as a single task. Let me explain why.

In order to start this task, you have to have some initialization; you have to have some overall process control which exists independent of that.

Secondly, if I have -- I don't know if this is being implemented as a PLC or if it is being implemented as an actual fully-implemented processor,

but if this is being implemented as a fully-implemented software system rather than using a PLC, then I'm probably going to have some kind of board support package with a number of low-level routines which are operating in memory at the same time, in order to get -- no?

MR. NGUYEN: No. For example, the Spin, which is the reactor protection system we use in the N4 series --

MR. HECHT: Yes.

MR. NGUYEN: -- which was built in the nineties, is completely custom-made. We require that you have the total source code of the system.

MR. HECHT: Well, you might still have the source code of that, but the point is that there is actual low-level routines, hardware interrupt surface routines.

MR. NGUYEN: No, no, no. No, no, no, there are no hardware interrupts. The only interrupts that occur are, I would say, the exceptions.

MR. HECHT: Timing --

MR. NGUYEN: No, no, not even that. But are the exceptions. For example, when you lose power, then there is an exception that is sent to the processor to say, well, you lose power in 5

milliseconds; do whatever you want to do, but in 5 milliseconds it is over.

MR. HECHT: And you also probably have another one servicing the reset switch.

MR. NGUYEN: Yes, that's right.

MEMBER BROWN: But external to the program cycle interrupt.

MR. HECHT: That is the point. So that is external to the program cycle. So it is a separate task.

MR. NGUYEN: No, no, no, it's not a task. It's a hardware signal that interrupts and that's a mechanism of the microprocessor that interrupts the execution, the current execution, leads to stop the execution at the specific address, and the specific address just says stop.

MR. HECHT: Okay. So this is being implemented directly as an interrupt --

MR. NGUYEN: That's right.

MR. HECHT: -- to the hardware?

MR. NGUYEN: That's right.

MR. HECHT: Okay. I guess there are many ways of implementing that, so that the failure mode -- what I have learned from what you have just described is, in this software architecture, failure modes would

be different than they are in another software architecture.

MR. NGUYEN: Oh, yes, and here --

MR. HECHT: Because what you are saying here is that you have not one, but probably countable, three or less, tasks which are running simultaneously on the processor.

MR. NGUYEN: No, no, no, no.

MR. HECHT: We would have to look at the detailed design --

MR. NGUYEN: Yes.

MR. HECHT: -- in order to say that, but --

MR. NGUYEN: In fact, it depends on the vendor. For example, in the Spin there is only one task. For other systems, it is less simple. Okay?

Here I am just giving the principles. Of course, when the principles are not completely adhered to, then you do have to do some analysis. That is what we currently do in my research center. It is to cope with real systems.

MR. HECHT: Okay. Because in a PLC, for example, there are actually many tasks.

MR. NGUYEN: Oh, yes. In a PLC that you buy from vendors who sell to the petrochemical

industry, it is usually multitasking software.

For 1E functions, most, at least those on which I'm really working, for which I have the source code, and I can verify at very low levels of detail, this is a single task.

MR. HECHT: But weren't Allen Bradley controllers, for example, used in diesel engine sequencings, start-up sequencing?

MR. NGUYEN: I admit I don't know the Allen Bradley.

MR. GEDDES: I think somebody did. I don't know --

MR. HECHT: Aren't those 1E systems?

MR. GEDDES: Yes.

MR. NGUYEN: I just wanted to say let's start with such a design. With such a design, you don't have -- can we go to the next slide?

MR. TOROK: Yes.

MR. NGUYEN: There are a number of, I would say, items that were in the list of 10 modes or mechanisms that are addressed by the watchdog. If one of the tasks or the single task crushes, for whatever reason -- it could be because of a division by zero. It could be because of a random single upset event that modifies a memory location, and that causes the

software to crash. It could be because of, well, incorrect code where you do have a division by zero, or whatever. The outputs will not be written within the framework of the 50 milliseconds.

Therefore, the watchdog timer will say something bad happened. I don't know what happened. I don't know what was the mechanism, but I will force the failure mode to be sent, the signal, saying do something.

So that covers, I would say, multiple possible mechanisms. The good designer is the one that is able to, I would say, cover as much as possible the possible failure mechanisms, so that that would lead to a known mode of behavior, not necessarily failure mode, but a no load of behavior.

And we can go through each of the items in the list, if you want. It is just to say that the notion of defensive measure is very closely related to the analysis of the possible failure mechanisms that we could have in a design.

MR. HECHT: Why one does this.

MR. NGUYEN: That's right. That's right.

However, in software, it is fairly different from the traditional way we analyze failure mechanisms. From the FMEAs I have read, the

traditional approach is to start with the components, small components, for which we have a list of known failure modes for these components.

Then we go up one level and try to see what are the effects of these failure, components failure modes, to some higher level of integration within the digital system, and so on.

We end up with, I would say, saying here are the possible failure modes of the digital system, of another system --

CHAIR APOSTOLAKIS: Can we come back now to the presentation? Are you satisfied?

MR. HECHT: Well, so far, I'm neither satisfied -- go ahead, George.

(Laughter.)

MEMBER BROWN: I will just make one observation. This process, I have no problem with this. I delivered about 30 or 40 systems designed with main operating loops of exactly this nature with exactly this architecture and feedback.

So Myron's right relative to there are other housekeeping functions that have to be performed which you find other methodologies to do that, so that you don't have anything interrupting that main constant cycle processing loop, which only does the

functions that are necessary to retrieve data, analyze it, determine whether you need to trip or not trip output, that kind of thing. That cycle just goes on and on and on and on.

The only time it stops is if you externally come in and say, "Stop. I want to change the data. I want to tell you to sample the test resistor" as opposed to the -- from a manual test standpoint.

So that thing runs all the time. It is a main operating loop. It is not interrupt-driven. I say that with a little bit -- because there's things called good interrupts and bad interrupts.

I had probably the smartest guy in the world explain this to me about 20 years ago, which I have probably forgotten all that. But stuff like, if you put data into buffers, you have to clear, you have to reset buffers when data is being converted and being placed in buffers. An external reset can clear those buffers because they don't interrupt the main -- if they don't clear the buffer, that's fine. You just get crappy data there the next time or you get no data, and the thing continues to run.

The watchdog timer is a form of an interrupt. In other words, it comes back and stops

everything, only it really stops it.

So that is the only point I am trying to make, is this process has been applied in roughly 100 operating reactors today.

MR. HECHT: You mean that architecture?

MEMBER BROWN: That architecture, yes. In fact, we said you will not design an interrupt-driven operating system of any kind. We said that, but we actually did, and it was so hard to make it work that we vowed we would never do that again.

MR. GEDDES: Or to make its behavior predictable.

MEMBER BROWN: Well, an interrupt-driven system is not determinate by nature. A main operator mode is determinate by nature. People will argue that it is predictable and repeatable.

An interrupt-driven system, you have to go and do it on a statistically-determinate basis, which is very, very hard to do, very hard to do. Even in a fixed main operating loop, you don't have a fixed time response. You may have a protection function of 250 milliseconds, for instance. So you go through five 50-millisecond cycles. You start some. You do some more. Because you don't want to spuriously trip stuff.

So what you end up doing, if you do a time response test on that, you have to figure, well, I may start, I may enter the time response test right after all the data is picked. So I have lost 50 milliseconds in my time or 49.99, whatever it is in that.

So you run a test once, and you get 200 milliseconds. You run it again, and you get 235. You run it again, and you get 220. So you have to do about hundreds of tests in order to get a consistent statistical basis to prove that you are really less than 250 all the time.

There's a process for doing this. It works very, very well as long as you stick with main operating loop, non-interrupt-driven systems, very important.

MR. GEDDES: And these are forms of defensive measures?

MEMBER BROWN: Oh, it is an extremely defensive measure, and it does blanket, it captures a lot of these issues. I mean I wasn't here when you did these little things, which are all very good failure mechanisms. Excuse me. I almost said the wrong word there. I don't want to do that again.

(Laughter.)

So, anyway, I will stop right there.

MR. NGUYEN: Can we go to the next one?

The only point I would like to say, there is one item that was very interesting in the list. It is task incorrect response. Because, of course, there is no, I would say, predefined defensive measure that will prevent all tasks from providing incorrect response. Then you need to have a very close look, and most of the time it is very application-dependent.

For other reasons, when we analyze digital systems and their software, we arrive at the conclusion that the main cause -- the part of the software that is most likely to cause failures of the digital system would be the applications.

There are defensive measures in what could be called the operating system that I would say, more or less, I would say, relieves the operating system from the accusation of causing failures. It will be mostly the application and in the OE. In fact, it is what we see when we look at the failures that affected the 1E systems and the non-1E systems. The main cause of that is software-related causes, are the applications.

Yes?

MEMBER STETKAR: Just go back a second.

You said the only interesting failure --

MR. NGUYEN: Oh, no.

MEMBER STETKAR: -- whatever you call it, is a task incorrect response. I am curious why the task early response is not interesting.

MR. NGUYEN: The task early response, if you think of a cyclic operation, at each cycle you have an answer.

MEMBER STETKAR: No, no. Well, it says, but in the worst case they constitute a spurious actuation.

MR. NGUYEN: Yes.

MEMBER STETKAR: Why is that of concern?

MR. NGUYEN: Oh, it is of concern, but, in fact, it is the same as task incorrect response.

MEMBER STETKAR: Okay. I mean, if you are taking that broad --

MR. NGUYEN: Yes, yes.

MEMBER STETKAR: Okay, fine. Thanks. Go on.

MEMBER BROWN: We actually had a self-test check for exactly that type of thing.

CHAIR APOSTOLAKIS: Well, is it time to take a break? Yes, it is.

(Laughter.)

MR. NGUYEN: I will finish in two minutes.

CHAIR APOSTOLAKIS: You will finish in two minutes?

MR. NGUYEN: Yes.

CHAIR APOSTOLAKIS: If they don't interrupt you. Okay, go ahead.

MR. NGUYEN: So our conclusion is that the application code is, I would say, probably the most dominant cause of software-related failures. We are speaking about determining values for PRAs. It is important to understand that. Because if you think that the dominant cause of software-related failure is the operating system, the software platform, then your beta factors will be very different.

If I say that I have two subsystems with different applications but the same platform, if it is the applications that are the dominant cause of software-related failure, it is not at all the same to say it is the platform which is the dominant cause. The beta factor would be very different.

In one case, I would say, if it is the platform, the beta factor will be probably not very far from one. And if I say it is --

CHAIR APOSTOLAKIS: You exaggerate.

MR. NGUYEN: Yes, I exaggerate. But I

guess for software-related failures, if it is the application, the beta factor might be much lower.

So the understanding of these defensive measures, the effectiveness of the defensive measures is important when you try to determine what are your values, your best estimated values, for your PRA model.

MR. HECHT: Okay. By defensive measures, in this case you mean the architecture?

MR. NGUYEN: The architecture, the design features, the design features in the software, for example.

MR. HECHT: Okay, the architecture and the design?

MR. NGUYEN: Yes.

MR. HECHT: You don't mean the process?

MR. NGUYEN: No.

MR. HECHT: Okay. Might I just make an observation that --

MR. NGUYEN: Yes.

MR. HECHT: -- you have defined in your particular case for the Spin controller?

MR. NGUYEN: Yes.

MR. HECHT: You called it, it seems to be a special case. There are other software

architectures that are in class 1E systems; there are non-class-1E systems that have different software architectures. So, obviously, I think we can agree that the failure modes obviously have to be tied to the architectures and to the design.

I also wanted to make the observation that, just as you were talking about, for hardware failure modes you start with the components.

MR. NGUYEN: Yes.

MR. HECHT: And my experience doing this work, the task is the equivalent of the component for software. So, if you have -- "if", and I'm not convinced, but I don't know your design -- if you have actually only one task --

MR. NGUYEN: Yes.

MR. HECHT: -- then there would be only one task there. It might have very few failure modes.

MR. NGUYEN: Okay, maybe we can continue that during the break.

MR. HECHT: But there are others.

MR. NGUYEN: Yes.

MR. HECHT: I guess that is the point.

MR. NGUYEN: But just my last point is that one of the four software-related events we had seen in the OE, the case where the self-test modes

were preventing the generation of the protection signal, it is typically a case where this rule of separate design, no interruption, and so on, was not put out. That was, I would say, one of the causes of the problem.

CHAIR APOSTOLAKIS: Okay, back at 2:55.

(Whereupon, the foregoing matter went off the record at 2:40 p.m. and resumed at 3:06 p.m.)

CHAIR APOSTOLAKIS: Back into session.

MR. TOROK: We would like to move on as quickly as we can into the discussion of that DAS report and risk insights, and so on. I think that we have pretty much made our points in regard to failure modes and effects, and so on. So I want to wrap that up very quickly here.

This is the next slide. You haven't seen this one yet, but the point is, so what are the CCF implications now that we have talked about mechanisms, modes, and effects, and so on?

As Thuy was explaining, mechanisms can be addressed to a large extent by defensive measures and/or diversity. So often, as Thuy explained, defensive measures can eliminate entire classes of failure mechanisms, which is a good thing. It also means that we can probably learn how to be more

efficient about our FMEAs, and, of course, we are going to work on that.

And diversity, while it can be helpful, isn't the only solution. We think what you are really looking for is to use the best aspects of both and be aware of those. But, if you are serious about protecting against common-cause failure, I think you have to be serious about looking at defensive measures.

CHAIR APOSTOLAKIS: So what do you mean by lines of defense? It may be more appropriate with different lines of defense?

MR. TOROK: That is a reference to the notion that, well, if you've got redundant trains that have the same functionality, typical of safety systems, right, then in a situation like that, diversity, platform diversity, doesn't really buy you much because the things you are most worried about are problems coming from the requirements or the application code, and diversity is not really going to help you there. Platform diversity is not going to help you.

As opposed to comparing two different lines of defense, where they have different functionality to start with, and typically, different

platforms to start with, right? So there you've got a lot more going for you in terms of using diversity as a protective mechanism for CCF. That is what that is about.

So what we would say is, when you are trying to figure out what combination of defensive measures and diversity attributes to use for a particular application, keep all that stuff in mind. Okay?

CHAIR APOSTOLAKIS: Now a place where the issue of diversity became real was adding a diverse shutdown system if the operator action was supposed to be --

MEMBER BROWN: Less than 30 --

CHAIR APOSTOLAKIS: -- less than 30 minutes.

MEMBER BROWN: Thirty minutes, right.

MR. TOROK: And that is exactly the case --

CHAIR APOSTOLAKIS: How would this supply to that?

MR. TOROK: That's exactly the case that Dave is going to explain. Okay?

CHAIR APOSTOLAKIS: Okay, good. Good.

MR. TOROK: We're here to help, you know.

So meanwhile, the next step, mechanisms, modes, and effects, what are the PRA implications? And the answer is just really that top line. It is the modes and effects you care about, not so much the mechanisms in PRA.

Now Dave is going to explain how that translated into what he did with that DAS example. Okay?

Now, still, you have to deal with this question of, what are the probabilities of failure? For that, it is true that understanding dominant failure mechanisms, and so on, may be helpful. But Dave is going to explain what was done in that particular evaluation for you. Okay?

Let's see. That's really all I was going to say about that one.

Now the bottom line here in terms of mechanisms, modes, and effects, here I am going to do this overstatement again.

"Failure modes in digital protection systems are well-understand." You know, that one is not so bad if we consider that the protection systems are relatively simple and we are talking about the modes and not mechanisms. So maybe that is not such an overstatement.

MEMBER STETKAR: Yes, when you say, "protection system", have you now focused your attention simply on tripping the reactor and not safeguards actuation systems?

MR. TOROK: A good question. For the purposes of this, I would suppose, being glib here, what I am talking about is typically systems where they are monitoring some parameters, comparing the values to a setpoint, and saying go or don't go.

CHAIR APOSTOLAKIS: So these safeguards included actuation.

MR. TOROK: So the answer is, yes, I would include ESFAS.

CHAIR APOSTOLAKIS: Actuation, yes.

MR. TOROK: Yes, and for those, typically, the system output is a one or a zero, and those are your only choices. For those, I think we have a pretty good handle on the failure modes. That is really all I am saying. Okay? And for the most part, they are not any different from what you have with the analog.

You can argue about mechanisms being --

MEMBER STETKAR: No, no, no.

MR. TOROK: Okay. Now we also noted that a lot of FMEAs are being done, maybe extensive FMEAs

on various pieces of equipment.

MEMBER STETKAR: Oh, wait. You jumped over that middle section there that said, "Common-cause effects are modeled in PRA for existing plants." I am curious to hear about that because I have looked at a lot of PRAs for existing plants that do not model common-cause effects of instrumentation and controllers. In particular, spurious actuations.

If we did that, the industry would not have spent Lord knows however many millions of dollars trying to integrate fire risk assessment into their wonderful internal events common-cause models, for example.

MR. TOROK: I am going to defer to Dave on that one.

MR. BLANCHARD: Actually, this bullet is mine.

(Laughter.)

And common-cause effects with respect to actuation of a system certainly are modeled in the PRAs. So there's quite a number of PRAs that have them.

MEMBER STETKAR: Failed to start, they are --

MR. BLANCHARD: Yes, they failed to

actually --

MEMBER STETKAR: I'm not interested in that one. That's the easy one.

MR. BLANCHARD: Yes, I understand.

MEMBER STETKAR: The one I am interested in is the spurious actuation of things when you don't want them to do that.

MR. BLANCHARD: And you are correct, there is extensive work going on right now to incorporate that in for NFPA 805, I think it is.

Yes, the spurious actuation leading to a transient event, I would say is incorporated fairly well, largely through initiating event.

MEMBER STETKAR: That's a surrogate for --

MR. BLANCHARD: Right, but during a transient, the spurious actuation, you're right, that largely is left out right now, and it is being added as a regular fire PRA.

MEMBER STETKAR: Thanks.

MR. TOROK: Very good. Okay.

Now the only other point that we were trying to make was that, while the failure mechanism may not be particularly important at the PRA level, they are very useful when you are evaluating digital systems and looking at the design stage to make sure

you have incorporated the right defensives against those things in terms of defensive measures. So an understanding of the failure mechanisms is still a very useful thing.

So what are we recommending here? Well, basically, we think that it is important to consider defensive measures in terms of evaluating your protection against common-cause failure and the adequacy of that protection. It would be nice to develop what we call here deterministic criteria for applying defensive measures.

Let's see. We are also looking at, and we think the work should continue here, looking at evaluating the defensive measures that are available, effectively estimating the coverage against failure mechanisms as a means to get a handle on reliability estimates for use in PRA.

Now the reason that is interesting is because, if you look at the data for failures of 1E systems, for example, and you are trying to create a statistical basis for doing something there, it is pretty difficult because there aren't a lot of demands typically, and these systems are designed to be exceptionally robust. So there aren't a lot of failures.

So, if you are trying to generate statistically-significant numbers, that might not be possible in any relevant timeframe. So we think it is useful to look at defensive measures this way.

I don't think it is unlike what is done in other parts of PRA, though, when you are talking about using expert elicitation, I guess, to estimate failure probabilities, and so on. So I don't think it is really out of line there.

And we think we ought to continue these efforts as part of the work we are going to do under the MOU and coordinate with NRC.

Okay, having said that, now we finally get to our final topic, which is risk insights and, in particular, this diverse actuation system analysis. There was a white paper that was submitted to NRC back in May last year, and the idea was we looked at -- it was a risk-informed look at the potential benefits and risks of an automated DAS that might be required per the ISG 2 of September 2007. At the time, of course, that was a burning issue in the NEI Working Group, which is why we were looking at it.

Now we published a final report on that late last year. We gave it to ACRS and the NRC in January this year.

Between the white paper and the final report, the conclusions didn't change. The methodologies didn't change, but we restructured it to make it easier to read. We moved a lot of the detailed number-crunching to the appendices, the idea being you can sit down and read through the whole report in one sitting now. So we tried to straighten it out that way.

Let's see. We also incorporated comments and tried to address comments from discussions with the NRC Task Working Group on this.

And we added one sensitivity study, which was actually suggested by NRC staff, and it had to do with the benefits, or relative benefits, of prevention versus mitigation as ways to address the common-cause failure problem.

Okay. Now, so in looking at the DAS, what we are going to talk about is this: first of all, think of this analysis as an example of how you can generate useful risk insights for digital systems using existing PRA methods. That is one thing we think it is useful for.

The DAS case itself was an example of how to do that. At the time, as I said, it was an interesting issue for the Working Group, and that is

why we were looking at it.

We did deterministic evaluations to figure out where you might need DAS. Now we are going back to the NRC policy statement and the SECY and how you deal with that.

Then we did probabilistic analyses to generate an assessment of the potential risks and benefits. Now to do that, we had to put in numbers. So we are going to talk about how we did that, where we got those numbers. Dave is going to answer the hard questions there.

We also had to somehow factor in this notion of failure modes and effects, along the lines of what we were talking about earlier: what do we care about? Do we care about task hang in the microprocessor? PRA doesn't care about that. Dave is going to tell you how that was handled in his analysis. Okay?

So we will show you what the results of the study were, including sensitivity studies that Dave performed and the impact of that on the risk insights that came out of it.

One thing that is really interesting, for me anyway it was interesting, that came out of Dave's analysis was, if you are going to put in one of these

DASs, these automated diverse actuation systems, it became obvious that there was this concern about spurious actuation of it, causing transients that, of course, have some risk associated with them, and you would rather avoid that.

But there are ways to do that, and the risk insights help you figure out some of the good design features that can help you with that. So Dave has got information on that.

And lastly, as I mentioned earlier, we did this thing with the understanding of the 30-minute criterion of the DAS a couple of years ago. There's now a new one. We can talk about the potential impact of the revised criterion, if we want to get into that.

So what are our key points? First, we think it is possible to generate useful risk insights right now using the tools that are available right now, even without precise knowledge of the failure mechanisms and the probabilities at the component level.

What we are trying to show here is you can do this if you keep your modeling level of detail appropriate for the application. Of course, we picked a particular application here. So this is a confined study. It showed that in this case the results were

insensitive to wide variations in assumptions, which is a good thing.

One of the reminders here is that we were looking at a particular individual system and saying, oh, man, what if there is a common-cause failure there? We are trying to treat it conservatively by adding a diverse backup and that kind of thing.

One of the things the analysis tells you is that picking out a component of a big, complicated system and trying to treat that conservatively doesn't always result in an overall result, what you were looking for. In other words, it may not improve the safety of the overall system. Maybe it could even degrade it.

That is one of the things that looking at it from a risk perspective brings, which I think is really valuable. The risk analysis, it doesn't look at the trees. It stands back and looks at the forest. Sometimes that is very valuable.

Right now, industry in various places is applying existing methods. I think the new plant guys are doing it or the vendors of the new plants are doing it, and applying PRA insights to help design their systems better, and in some cases operating plants are doing it as well.

Now this last note here refers to some ACRS statements expressing basically skepticism as to what you can and cannot do with risk methods right now, how far we can go with this.

Some of those statements have been construed basically to mean that it is not possible to generate risk and be concise. Now, basically, it is not possible to do what we think we did. So we are thinking that it would be interesting to revisit that question later, after we have gone through and shown you what was done and given you a chance to comment on that. Okay?

Any questions?

(No response.)

Thank you.

Dave, please, help.

MR. BLANCHARD: All right. The starting point for this analysis turns out to be Branch Technical Position 19. Branch Technical Position 19 required to analyze each design basis event, assuming a coincident for a common-cause failure in the reactor trip system or the ESFAS.

We also have some additional guidance where the staff expressed the desire to limit credit for operator action, should we need it, and

demonstrating adequate defense in-depth against common-cause failures to timeframes greater than 30 minutes. And if we need this operator action in timeframes less than 30 minutes, they would like to see an automated diverse actuation system. That, of course, is from ISG 2 of the D3 Task Work Group. It is sometimes referred to as the HOV lane for licensing review of digital upgrades and I guess digital systems for the new plants.

Now that ISG has changed, as Ray mentioned. Since we started this analysis, it has been modified to reference ISG 5 in the area of credit for operator actions. It appears that there is an alternative at this time to the 30-minute criterion where we can go in and do human factors engineering analysis, should we want to credit operator actions in less than 30 minutes. That does provide us some needed flexibility. However, it doesn't address all of the accidents.

The more rapidly-evolving events, like the large LOCA and large steam line breaks, may not be helped by that particular ISG. We can get into the reasons for that a little later, if you would like.

So, at any rate, we still think this analysis of the risks and benefits of the automated

DAS system has some relevance here, even with that ISG.

I guess what our objective here is to demonstrate that we can, in fact, use existing risk techniques in order to evaluate systems such as this in PRA in a digital INC context.

All right, next slide.

All right. Ray briefly outlined the approach we took in doing this analysis. We first began with a set of deterministic analyses, and then moved on to an accident sequence analysis. I would like to talk a little bit about the deterministic analyses that were performed.

Our purpose or our objective in performing these deterministic analyses is to identify precisely just which transient and accident sequences need an automated DAS, as described by Branch Technical Position 19 and the ISG.

For the purpose of doing this, we had four plants volunteer some of their thermal hydraulic models and also their PRAs in order to assess what accidents and transients would fall in the category of needing operator actions.

CHAIR APOSTOLAKIS: That sounds like a pretty large effort.

MR. BLANCHARD: A pretty what?

CHAIR APOSTOLAKIS: A significant effort.

MR. BLANCHARD: Actually, this is a relatively simple application.

CHAIR APOSTOLAKIS: But relative to what? I mean you had the utilities involved. Presumably, your time was significant. So all this because of the 30 minutes?

MR. BLANCHARD: All of this because it seemed we needed an automated DAS in order to license --

CHAIR APOSTOLAKIS: It would have been much more expensive.

MR. BLANCHARD: -- and other things as well, right, which we will touch on here for this analysis. Right.

At any rate, we had both the PRAs as well as the thermal hydraulic models for these four plants. We had a Westinghouse 2-loop plant, a Combustion Engineering plant, a Babcock and Wilcox plant, and a BWR 3. So we touched on each of the four major reactor vendor designs in the U.S. as a part of this evaluation.

The scope of the evaluations include the full spectrum from the internal events PRAs for each

of these plants. The transients, the full spectrum LOCAs, steam line breaks, and ATWS as well.

Now we had the thermal hydraulic analysis from the PRAs. So, with respect to the transients and the ATWS, we were pretty much able to rely on what they already had in the PRAs to assess the relevance of transients and ATWS to the 30-minute criterion.

For transients, we are pretty much losing inventory at decay heat rates. The loss of inventory is relatively slow. Timeframes are beyond 30 minutes for most of the transients, and therefore, we came to the conclusion we really didn't need the automated DAS to meet the ISG 2 for the transients.

MEMBER STETKAR: You don't need it for a B&W plant on a loss of feedwater transient?

MR. BLANCHARD: No, actually, that was longer than --

MEMBER STETKAR: You mean B&W steam generators dry out in more than 30 minutes?

MR. BLANCHARD: Well, they dry out. They dry out quickly, but then you move to once-through cooling or feed-and-bleed. Okay? At that point, you are in or beyond 30 minutes there. You may be into feed-and-bleed, but you certainly aren't into the point where you're getting any fuel damage within 30

minutes.

We did look at some additional transients or additional failures on top of the transients, such as stuck-open safety valves on top of the transient-initiating events by itself. That's nothing that is required by BTP 19. You don't have to assume additional failures beyond the transient and the common-cause failure. We did that just to get the timing in the events down, so that we could do the sensitivity studies.

For ATWS events, what we found for the PWRs, well, what we found for all plants is we already have an ATWS system for each of the plants that is there. It is diverse to the reactor trip system, and it is there to cope with ATWS events. So our assessment was largely to decide whether or not we also needed some automation of ESFAS during the ATWS events.

For PWRs, we discovered we would not get a safety injection signal for ATWS conditions. And even if we did, quite often, the reactor pressure would be high enough that we would not have the ability to inject to the reactor with the safety injection system. So we concluded we didn't need diverse actuation for ESFAS during ATWS for PWRs.

For BWRs, you can get an ECCS signal during ATWS to start your injection systems. But the EOPs are written to defeat the ECCS during ATWS, and the benefits of that are you lower power; you lower loads on the containment; you buy time for injection to SLC.

So the issue there is, if we provided a diverse actuation system during an ATWS for a BWR, we would have just kind of created an additional system that they would have to defeat in order to implement their EOPs. So we decided the diverse actuation system was not necessary for ATWS conditions.

That leaves the LOCAs and the steam line breaks. For those, we did review the existing thermal hydraulic analysis from these four PRAs, but, in fact, we found we needed to do some additional thermal hydraulic analysis.

So, if we could go to the next slide, we will start off with the loss of coolant accidents. We found each of these plants had additional definitions of what they meant by large, medium, and small LOCAs. So we came up with a consistent definition for at least the large LOCA category for these events.

We decided to call the large LOCA, any LOCA size that low pressure injection would be

effective in providing adequate cooling. So our thermal hydraulic analysis for the PWRs was directed at deciding what that break range was.

For Westinghouse plants, any break greater than 4 inches, for the Westinghouse 2-loop plant, any break greater than 4 inches, we found their low pressure injection system could provide adequate core cooling. We found the same number for the Combustion Engineering plant, and a little bit larger break for the B&W plant.

Now, at this point, we are going to define large LOCA or redefine the large LOCA spectrum as being any break larger than these for these plants. All right?

What we did in a second set of thermal hydraulic analyses at this point is we did a test to see if the low pressure injection system did not actuate for breaks at each of these sizes, how long it would take to get to core damage, in fact, if no injection systems worked at these break sites.

What we found was for the Westinghouse plant at 4 inches we had two hours before we had any fuel damage. The CE plant was four hours, and the B&W plant was 45 minutes.

Now the differences in these numbers has

to do with the volume of the primary coolant system, the size and the pressure of the accumulators. That is the reasons for some of these differences.

But, with respect to our diverse actuation system risk analysis, what this meant was the only system that we needed to actuate automatically with a diverse actuation system was the low pressure injection system because we had more than 30 minutes in our defined break spectrum here for the entire -- well, for --

MEMBER BROWN: And that was for BWRs? I mean I'm reading your table all the way down. So, I mean, less than 15 minutes is the category you are talking about then?

MR. BLANCHARD: No.

MEMBER BROWN: For the low pressure injection system.

MR. BLANCHARD: This is the PWR. I'm just talking about the PWR category.

MEMBER BROWN: You said you needed a DAS for that. Yet, it is greater than 45 minutes.

MR. BLANCHARD: No.

MEMBER BROWN: The first one is.

MR. BLANCHARD: If we are going to provide a DAS, it would only be needed for low pressure

injection. Okay? That is because for breaks for which you need high pressure injection you have longer than 30 minutes before you would need to start high pressure injection.

Obviously, I'm not --

MEMBER BROWN: I haven't connected the dots here. I'm just reading the thing.

MR. BLANCHARD: I understand.

MEMBER BROWN: It says, "purpose: mitigated by low pressure injection. Results: it doesn't work. It doesn't matter."

MR. BLANCHARD: Well, it does matter.

MEMBER BROWN: You've got greater than two hours.

MR. BLANCHARD: Right. I do need to get low pressure injection systems running. Those have much longer than 30 minutes to do that.

MEMBER BROWN: Yes, but you could do it manually is the point.

MR. BLANCHARD: And the only thing you need is low pressure injection.

MEMBER BROWN: But you can do it manually.

MR. BLANCHARD: Oh, I'm sorry. I can't do it manually for the double-ended guillotine rupture, right? It is much shorter than two hours or 30

minutes even for the double-ended guillotine rupture. It is only going to be three or four minutes.

CHAIR APOSTOLAKIS: What you say there, Dave, is that tying to core damage without low pressure injection, for Westinghouse, it is two hours; for CE, it is four hours --

MR. BLANCHARD: At a break size of 4 inches and the smallest large LOCA. For the largest small LOCA, it is just a few minutes.

CHAIR APOSTOLAKIS: Okay.

MR. BLANCHARD: And it is that largest small LOCA that I need the diverse actuation system for.

MEMBER BROWN: What is the largest small LOCA?

MR. BLANCHARD: Double-ended guillotine break in the largest pipe.

Oh, I'm sorry. Did I say small LOCA?

MEMBER BROWN: Yes.

MR. BLANCHARD: The smallest end of the large LOCA range is 4 inches. The largest --

MEMBER BROWN: Okay. It's not in your table.

CHAIR APOSTOLAKIS: It's not in the table. That's what is confusing.

MEMBER BROWN: That's my problem.

MR. BLANCHARD: I'm sorry. Okay.

CHAIR APOSTOLAKIS: You should have had a "C" that says, for a large LOCA break, you have a few minutes.

MR. BLANCHARD: Yes. That is not in the table.

MEMBER BROWN: You should have had another line up here that said, "Large LOCA, double-ended guillotine". I'm saying you never need a DAS or anything. Why bother? Who cares?

MR. BLANCHARD: My apologies.

CHAIR APOSTOLAKIS: You only want to put good stuff on the slides, right?

(Laughter.)

MR. BLANCHARD: This is good stuff. It just wasn't enough good stuff.

Yes, for the largest break, we only have a few minutes.

CHAIR APOSTOLAKIS: Okay.

MR. BLANCHARD: And that's the break that decides you need to automate low pressure injection.

MEMBER BLEY: The 4 inches, 4 inches and larger, all you need is low pressure injection? That's why they did that?

MR. BLANCHARD: That's the conclusion. That's right.

MEMBER BLEY: If it is at least 4 inches or bigger, you only need one system to take care of that?

MEMBER BROWN: Isn't a double-ended break larger than 4.5 inches?

MEMBER BLEY: The first step, they look for a LOCA that one system will take care of. That was the bigger than 4 inches with low pressure injection. Out of all of those, it is the double-ended guillotine that sets the shortest time.

MR. BLANCHARD: Right.

CHAIR APOSTOLAKIS: So up there, then, under Westinghouse, maybe you should have said between 4 inches and 6 inches?

MEMBER BLEY: No, four.

MR. BLANCHARD: The analysis that generated the two-hour number is for the 4-inch break.

CHAIR APOSTOLAKIS: Right.

MR. BLANCHARD: Yes.

MEMBER BROWN: So that is greater than 4.5 inches in diameter?

MR. BLANCHARD: Can be handled by low pressure injection.

MEMBER BROWN: In the coolant loop there is a pipe greater than 4 inches in diameter?

MR. BLANCHARD: Oh, this is the break size. This is the effective break diameter here. This could be a 4-inch hole in that big pipe or it could be --

MEMBER BROWN: It could be a double-ended break that is greater than 4.5 inches also.

MR. BLANCHARD: Right, but low pressure injection systems will handle that break.

MEMBER BROWN: Won't.

MR. BLANCHARD: Will.

MEMBER BROWN: A double-ended break?

MR. BLANCHARD: Yes.

MR. SIEBER: Yes, because the pressure goes down so fast.

MR. BLANCHARD: Because the pressure goes down --

MEMBER BLEY: By design.

MR. BLANCHARD: Right, by design.

MEMBER BROWN: Now I understand it now that you have explained it to me. Other than that, I was totally lost. When I read this, I said, why are we bothering with this report?

MR. BLANCHARD: And it is because there

are bigger breaks than this that are much faster.

CHAIR APOSTOLAKIS: But that inequality there, Westinghouse greater than 4 inches, should have been really approximately equal to?

MR. BLANCHARD: Yes.

CHAIR APOSTOLAKIS: Then the rest of it applies?

MR. BLANCHARD: Yes.

CHAIR APOSTOLAKIS: Because if you look at it now, you will say, well, okay, if it's a large break, then this doesn't apply.

MR. BLANCHARD: Yes.

CHAIR APOSTOLAKIS: So it should have been approximately.

MR. BLANCHARD: Yes.

CHAIR APOSTOLAKIS: Okay, good.

MR. BLANCHARD: Okay. Now let's take a look at the BWR. We did the same type of analysis for the BWR, and we found that the smallest end of the large break spectrum was around 4.8 inches. Okay?

Then we did the second analysis, if we have no injection at all. For that break size, approximately 4.8 inches, how long does it take to get to fuel damage? And it turns out to be only around 15 minutes. All right?

Now the first question you might ask, how come the BWR is so different than the PWRs in terms of the response to these break sizes? And the answer is the PWRs have accumulators that are injecting during the large breaks that are extending the time to fuel damage that the BWRs don't have. Okay?

MR. SIEBER: Plus, they have loops.

MR. BLANCHARD: I'm sorry?

MR. SIEBER: They have loops, too.

MR. BLANCHARD: Oh, yes, on the steam generators and everything. Yes. Okay.

Now the BWR for this analysis, we made the assumption that condensated feedwater was still available. So we are pumping the hot oil in. In addition --

MEMBER STETKAR: I noticed that in most of your cases. Why did you make that assumption? Why is that?

MR. BLANCHARD: Why did we make the assumption? First, from a BTP-19 perspective, we can make the assumption that we have --

MEMBER STETKAR: But the rules determine that it is available for plants --

MR. BLANCHARD: The rules allow us to do that. We are trying to determine for what accidents

we need this diverse actuation system for. So we accredited the condensate system here because we could.

MEMBER STETKAR: So you accredited the non-safety systems here because you could?

MR. BLANCHARD: Because we could. Still, it was only 15 minutes.

MEMBER STETKAR: In this particular case. In other cases, they bought you more than enough time.

MR. BLANCHARD: Bought us more than enough time in other cases, that is correct.

So what is happening here is that, for the entire large break spectrum, the entire large break spectrum is less than 15 minutes. It is less than 30 minutes in the ISG. What that means is that we have to automate, the DAS would have to automate both high and low pressure systems in order to meet the ISG.

So it is the difference between the BWRs and PWRs.

MEMBER BLEY: Just following up what John said, we introduced this as using PRA to find the scenarios that matter. The PRA would have scenarios that don't have those systems working, which meant some other situations might have been cases that none

of the facilities --

MEMBER STETKAR: Well, and my safety-related software might have shut off the feedwater condensate systems, for example, because it decided that I had something else going on that needed those shut off.

MR. BLANCHARD: If it had something in --

MEMBER STETKAR: Like a five-level signal in the vessel, which would shut off anything.

MR. BLANCHARD: Not condensate. Not condensate. It doesn't shut off condensate.

MEMBER STETKAR: It isolates feedwater. A level 9 signal closes the main feedwater isolation valves. I can't get condensate through those valves very easily.

MR. BLANCHARD: Okay, not at all plants.

MEMBER STETKAR: Okay. Not at all --

MR. BLANCHARD: Not at this plant.

MEMBER STETKAR: okay, maybe not at this plant, but at a lot of them it does.

MR. BLANCHARD: Right. Okay. All right, not at this plant though. All right.

So that is the additional thermal hydraulic analyses for LOCAs that we had to do for the LOCAs.

Now if we can move to the next slide, we also looked at steam line breaks. Again, we will do the PWRs and then the BWRs.

For steam line breaks, what we are talking in the PWRs is blowing down a steam generator and then looking at the plant response to that blowdown. We ran four cases.

We ran a case where ESFAS was successful, just to get baseline. By ESFAS here, I mean a couple of things.

One of them is making sure that feedwater isolates and MSIVs go closed. That is a secondary side of the plant part of ESFAS. The other is, when you get your safety injection signal, does safety injection start? So there's different ESFASs that we are talking about here.

We assumed in the first case both were successful. Then we ran three additional cases as sensitivity studies.

The first case was to see what the primary system conditions were if no safety injection actuated, but main steam isolation and main feedwater isolation did.

Case C here is what would happen if we didn't get steam line isolation or feedwater

isolation, but we did have some reduction.

The final case was, what if we had neither safety injection or feedwater isolation and main steam isolation?

What we found when we ran these cases is that each time we disabled one of the ESFAS functions, the fuel and the primary coolant system conditions were actually better, more benign than if these ESFAS systems worked. The primary coolant system pressures and temperatures were lower. Fuel temperatures were lower.

MEMBER STETKAR: I was curious, since we have until two o'clock in the morning for this -- (laughter) -- your second case ran that steam line break without an SI signal.

MR. BLANCHARD: Yes.

MEMBER STETKAR: And the notes say that the pressurizer repressurizes with reflood and water flow from the PORV at 1.4 hours. How does it repressurize if I don't have an injection?

MR. BLANCHARD: No, no, no.

MEMBER STETKAR: Or are you assuming I did injection at 30 minutes because, by definition, the operator starts --

MR. BLANCHARD: Okay. I will have to

examine what you are looking at here. The second case, you shouldn't be getting the PORV --

MEMBER STETKAR: I'm looking at the table in the report, Table C-1, Case 2B, which says, "PWR main steam line break without" -- okay, it's page C-12 in the EPRI report.

MR. BLANCHARD: Yes, we can pull that up because --

MEMBER STETKAR: Case 2B says it's the same as 2A, which is main steam line break, without SI. The comments says, "Pressurizer repressurizes with reflood and water flow from the PORV at 1.4 hours."

MR. BLANCHARD: I think that is an excellent question.

(Laughter.)

MEMBER STETKAR: But this one says it repressurizes in 23 minutes, but that's with the SI. I figured out how that got there.

MR. BLANCHARD: We will have to go back and look at this case, but we may have assumed charging was still --

MEMBER STETKAR: C-12. I couldn't figure out whether you were assuming the operators start injection manually at 30 minutes, and that's how I was

getting --

MR. BLANCHARD: I suspect that we did not trip charging. And of course, it wouldn't receive a signal to trip.

MEMBER STETKAR: As long as my really smart system doesn't know that it wants to shut off the charger.

MR. BLANCHARD: Right.

CHAIR APOSTOLAKIS: Do you have the PDF page number?

MEMBER STETKAR: The PDF page number, this is EPRI Report 1016721.

CHAIR APOSTOLAKIS: The cost/benefit.

MR. BLANCHARD: There's C-25, right?

MEMBER STETKAR: Page 86 on the PDF file.

CHAIR APOSTOLAKIS: Did you go to page 86?

MR. BLANCHARD: Yes.

MEMBER STETKAR: Here you go. Case 2B. Just scroll up or down, whichever -- there you go. Case 2B. Now scroll over a little bit.

Over in the comments section, on the 2A case, it says it refloods and pressurizes at 23 minutes, which I can understand that. High pressure injection is on.

And you get reflood and repressurizing at

1.4 hours.

MR. BLANCHARD: Right. Right. I suspect they left charging out.

MEMBER STETKAR: Okay.

MR. BLANCHARD: Yes, and that's refilling and charging as opposed to a large volume with high pressure safety injection, and that is the difference for the time.

MEMBER STETKAR: Okay. This SI signal doesn't isolate charging then? Okay.

MR. BLANCHARD: The SI signal would probably bring charging in --

MEMBER STETKAR: It depends on the plant design.

MR. BLANCHARD: Yes, it does. Yes.

MEMBER STETKAR: Okay.

MR. BLANCHARD: Okay?

But, at any rate, what is happening here, the cases where we're not actuating as fast are actually more benign than the case where ESFAS actuates. What is happening is that a steam generator is blowing down.

The primary system is cooling very rapidly. This is the largest steam line break, a double-ended guillotine rupture of the steam line.

MEMBER STETKAR: Okay. Did you take credit for the MSIVs closing on this one?

MR. BLANCHARD: For this, for the case without SI, we took credit for the MSIVs closing. So this one --

MEMBER STETKAR: So you're only blowing down a single steam generator.

MR. BLANCHARD: Right. Blowing down a single steam generator. And what happens is that, during the rapid cooldown, SI comes on, fills the primary systems. You dry out the steam generator because you have isolated feedwater. The primary system starts heating up on decay heat. Because you have filled it with water, as it heats up you hydro the primary system and lift the relief valves.

MEMBER STETKAR: If the MSIVs don't close, does this behave much differently?

MR. BLANCHARD: If the MSIVs stay open, then what happens is that the pressure in the primary system stays low because whichever steam generators you're making up to, they are going to be depressurized because the MSIVs didn't go closed. Right? And the primary coolant system pressure will stay low, and you will be removing heat from the steam generators at near atmospheric conditions.

So primary system conditions are actually less severe for the cases where the ESFAS doesn't work than cases where it does. So, regardless of the timing for this event, we made the assumption I don't think we need ESFAS in PWRs for the main steam isolation or feedwater isolation, nor do we need it for safety injection for the steam line breaks.

MEMBER STETKAR: As long as all of the plants have enough charging to make up for the primary --

MEMBER BLEY: For this particular plant.

MEMBER STETKAR: -- for this particular one plant.

MR. BLANCHARD: That's a good point. One thing we checked as part of this transient was, what does the level get to during the blowdown due to shrinkage? And it doesn't get anywhere near the top of the fuel load.

So what you need charging for is to take care of any leakage that may be occurring from the primary coolant system subsequent to the blowdown, and that takes a long time before you would ever get to the top of the core. That's much, much longer than 30 minutes. So you still wouldn't trip.

MEMBER BROWN: So you exhaust the

pressurizer?

MR. BLANCHARD: Yes, the pressurizer is empty for this.

MEMBER BROWN: So you drain it due to the cooldown?

MR. BLANCHARD: Due to the cooldown.

MEMBER BROWN: That's a lot of cooling.

MR. BLANCHARD: Yes.

MEMBER BROWN: I am not used to hearing drawing bubbles in the reactor vessel; that's all.

MR. BLANCHARD: Right.

MEMBER BROWN: It's not part of my background.

MR. BLANCHARD: Oh, I suspect not.

All right. Feedline breaks, PWRs --

CHAIR APOSTOLAKIS: Wait. There's a question.

MR. BLANCHARD: Oh, I'm sorry.

MR. WATERMAN: This is Mike Waterman, RES.

What was the period in the fuel cycle at which you did the large break? I was just wondering about reactor recriticality when you do your overcooling. If you are at the end of cycle, you really have no boron shim in the plant. So, if you overcool it, you put a lot of positive reactivity in

there.

That is the purpose of ESFAS, part of it is to borate the core.

MR. BLANCHARD: Right. We ran this, I believe, toward the end of the cycle to account for that. Now remember that we have a successful reactor trip during this particular event. So all the rods have gone back in, and there may be a small return to power. But as the primary system heats back up due to the decay heat and the small return to power, then it terminates that before you ever get to it.

MEMBER BROWN: It self-terminates it.

MR. BLANCHARD: Right. Now it may be different if you have a stuck rod on top of it.

Now, for the BWR, we also looked at steam line breaks. Now this is going to be steam line break outside the containment. Our ESFAS here is MSIV closure as well as actuation of the ECCS.

Again, for this particular event, we made the assumption that condensate was still available. What happens during this event is the reactor coolant system depressurizes through the steam line. First, I have the steam line break, and the MSIVs don't go closed because the ESFAS didn't work.

Now if you believe you can actually have

BWR main steam isolation valves that will remain open during this blowdown, then that is what we assumed here.

I got comments back, people said, during that blowdown, those MSIVs, you know, just the flow will draw those MSIVs closed, but we didn't credit that here. Okay?

So we blow the reactor down through the main steam line outside the containment, and again, the condensate system was available to pump water into the primary coolant system. With that additional inventory available during this event, we had some three hours before we needed to actuate any additional --

MEMBER STETKAR: Ray, can you toggle back to the table, please, in the report?

MR. TOROK: You betcha.

MEMBER STETKAR: And go down to the next page where you actually have the PWR steam line breaks.

On Case 4B, which is inadvertent relief valve opening without ECCS --

MR. BLANCHARD: Yes.

MEMBER STETKAR: It says core damage in 27 minutes. Why do I not need DAS for that?

MR. BLANCHARD: IORV with reactor trip.

MEMBER STETKAR: I mean, if 30.0000 minutes is defined as the difference between purely black and purely white, this would seem to apply to the purely black side of that line.

MR. BLANCHARD: This is a case without the ECCS or without the condensate system. Okay? This is one of those sensitivity studies where we are taking a look at some failures on top of the transient, in addition to the common-cause failure of the ESFAS.

So, in order to get this timeframe --

MEMBER STETKAR: This is not an initiating event of an inadvertent opening of a relief valve?

MR. BLANCHARD: The first case, Case 4A, is an initiating event. The inadvertent relief valves turn out and the BWR doesn't trip the reactor. It remains at power with this inadvertent relief valve, and where you get the reactor trip is as the steam flow, you know, the steam flow and the feedwater flow, you know, you will have a mismatch at this point, but the feedwater flow will pick up.

But the plant will stay at power, and what you are doing --

MEMBER STETKAR: It depends on how big the opening is. You might go out on high power.

MR. BLANCHARD: Actually, yes, it depends --

MEMBER STETKAR: You might go out on 110 percent power, depending on --

MR. BLANCHARD: You could, but what we found here was that we were staying at power and heating the suppression pool, and it took 27 -- no, it took -- yes, it took 27 minutes. You know, that 27 minutes in the table might be a typo. You notice it takes 27 minutes to get to the high drywell pressure by heating up the suppression pool.

MEMBER STETKAR: Yes, I noticed that on the first one, but I don't have the --

MR. BLANCHARD: I have that same 27 minutes under Case 4B, and I think I need to -- that is too much of a coincidence to me. I need to go check that.

MEMBER STETKAR: But the delta from 4B to 5A, for example, at the time of core damage to the vessel for each is comparable.

MR. BLANCHARD: That's seems reasonable.

MEMBER STETKAR: I mean I can't tell the difference. If that is supposed to be 37 minutes versus 27 minutes, you certainly can't tell that.

MR. BLANCHARD: I will need to go back and

look at Case 4B.

MEMBER STETKAR: Well, the more interesting case -- I had a question on 4B, but the more interesting case is 5A.

MR. BLANCHARD: Yes.

MEMBER STETKAR: That's 11 minutes.

MR. BLANCHARD: Right. What we are assuming here is a loss of feedwater, complete loss of feedwater. Now we are not crediting condensate here.

MEMBER STETKAR: That's right. This is feedwater isolation, for example.

MR. BLANCHARD: Yes, and we get the reactor trip on the loss of feedwater and we are imposing the stuck-open -- oh, I'm sorry.

MEMBER STETKAR: No, this is a steam line break.

MR. BLANCHARD: A steam line break. Excuse me. Yes. Right, this is the steam line break without condensate.

MEMBER STETKAR: Without condensate.

MR. BLANCHARD: Yes.

MEMBER STETKAR: And that's 11 minutes?

MR. BLANCHARD: Yes.

MEMBER STETKAR: But the generic conclusion is that I don't need DAS for any BWR steam

line break.

MR. BLANCHARD: We are looking at scenarios here to get timing, including those beyond those required to be analyzed.

MEMBER STETKAR: Oh, okay.

MR. BLANCHARD: All right?

MEMBER STETKAR: I got it. This one is beyond Branch Technical --

MR. BLANCHARD: Beyond BTP-19, but that doesn't mean we're not looking at it. Okay?

MEMBER STETKAR: Okay. Thanks. Thanks.

MR. KURITZKY: Excuse me one second. Alan Kuritzky from Office of Research.

Just to follow up on Dr. Stetkar's comment, I understand that because it is not called for in BTP-19, or whatever, that you don't analyze it, but you are doing a risk analysis. So you need to consider all the contributors to the risk analysis.

MR. BLANCHARD: Yes.

MR. KURITZKY: So the fact that you have a case that can potentially result in core damage in 11 minutes, I think you are probably premature to rule that out for all BWRs as being something that needs to be considered in that calculation.

MR. BLANCHARD: And when we get to the

accident sequence quantification itself, we include a much broader spectrum of accident sequences than are just --

MR. KURITZKY: So it gets folded back in?

MR. BLANCHARD: Yes.

MR. KURITZKY: Thank you.

MR. BLANCHARD: All right. So for the purpose of figuring out which accident sequences, in accordance with BTP-19 and ISG 2, require the automated DAS, these are the additional thermal hydraulic analyses we did.

We can go to the next slide.

So, from these analyses, we determined we really didn't need it for transients. We didn't need it for ATWS. We didn't need it for the steam line breaks. We were left the LOCAS, and it was the large portion of the LOCA spectrum we might need the DAS for.

For the PWRs, all we needed to automate was low pressure injection systems. For the BWRs, we found that we would have to automate both high and low pressure injection.

MEMBER BLEY: And this look doesn't consider that the DAS would be good because it reduces the risk. This only is you need the DAS to meet the

Branch Technical Position --

MR. BLANCHARD: So far, that is all the farther we have gone. Which accident sequences would we have to have and ask for in order to meet the ISG? Yes, that is all the farther we have gone. Okay.

Now, as we go along, we will broaden the scope of all the accident sequences and find out where it has benefit. Then we will do sensitivity studies on that, to try to expand the sequences which aren't considered by Branch Technical Position 19.

All right, we have decided what needs to be actuated and for what accident sequences. The last step of the deterministic process is to decide how the DAS should actuate these systems. This was actually an iterative process. We would do the risk analysis under an assumption as to how the DAS was actuated, and then we would modify that, some insights that we got out of the analysis.

Where we ended up is actually what we have listed here on this slide. We elected to have multiple diverse indications that were clearly indicative of the accident sequence for which the DAS was required before the DAS actuated these systems.

By multiple diverse indications, we are saying for PWRs we would like to have both a low

pressurizer pressure signal and a high containment pressure signal before the DAS actuated.

For the BWRs, a low reactor level and the high drywell pressure before it actuated. Now why these multiple and diverse signals, and it is to address the potential for spurious actuation of the system. We will see the quantitative reasons for that in a few slides. Okay?

Another feature of this diverse actuation system is we decided to require instrument AC for it to actuate. We didn't want a loss of an instrument AC bus, you know, shutting this thing off, again, for spurious actuation purposes.

Then we wanted multiple trains of this system to initiate the actuation, not just a single train. That was so that a single failure couldn't cause spurious actuation.

Again, this was iterative, and we ended up here as a part of the analyses that we did.

Next slide.

All right, that kind of sums up the deterministic analysis we did. Now we are going to get into the accident sequence quantification itself.

We will first evaluate the potential benefits of the automated DAS, and we will measure

that in terms of what it does to reduce the core damage frequency in the presence of a digital common-cause failure.

We will also look at the reduction in the release frequency and offsite consequences associated with sequences for which the DAS might be implemented, and then we will do a value impact analysis.

We will also take a look at the potential risks associated with the automated DAS. Again, these risks are those resulting from its potential for spurious actuation.

What we want to do here is just make sure that those potential risks are less than the benefits that we are getting out of the automated DAS. For the purpose of doing this evaluation, we had 10 plants volunteer that are PRAs. We had five PWRs, a Westinghouse 2-loop plant, a Westinghouse 4-loop plant, two CE plants, a B&W plant, and a BWR 2, 3, 4, 5, and 6. We had 10 plants all together.

When we got done with the base case evaluation for this, we took a look at the results and documented the reasons for the results and converted them into deterministic risk insights. Then we had made a number of assumptions, which we will talk about as a part of going through the accident sequence

analysis.

We did sensitivity studies and uncertainty analysis around a number of the more important assumptions to see how it influenced the results.

Next slide.

Okay. This slide has some 56 numbers on it, and I promise not to go through all of them.

CHAIR APOSTOLAKIS: How about the numbers on the upper righthand side?

MR. BLANCHARD: Yes. That's the one we want to focus on?

CHAIR APOSTOLAKIS: Yes.

MR. BLANCHARD: Okay. We can talk about some of the others. The LOCA-initiating event frequency comes from 1829, SECY NUREG-1829. It is a generic source of data that we use.

I also want to note on the left side of -- well, first of all, the top half of this table is the quantification of the benefits. The lower half of the table is quantification of the risks.

We will talk about the benefits first. On the upper lefthand side of the table, you will see the events for which the benefits apply. You will see that we quantified accident sequences more than just the large LOCA. Okay?

MEMBER STETKAR: Dave, on your slide, this is probably not relevant. I just want to understand.

The lower lefthand corner, you just skipped over. It said, "Spurious MSIV closure is 2.4 times 10 to the minus 3 per year."

MR. BLANCHARD: I'll get the basis for that.

MEMBER STETKAR: I don't care about that one.

Spurious reactor trip is also 2.4 times 10 to the minus 3 per year?

MR. BLANCHARD: Yes, there's a reason for that. I have a slide that explains it.

MEMBER STETKAR: All right. Thanks.

MR. BLANCHARD: All right.

The events we considered that would benefit from the proposed automated DAS are more than just the large LOCA. Once we provide an automated DAS for, in this case it is a BWR, the BWR results that we are looking at for the BWR -- remember, we are automating both high pressure and low pressure injection.

What that means is that more than just the large LOCA is going to benefit from the automated DAS. The full spectrum of breaks are going to benefit from

automated DAS, including the small LOCAs, even those events for which there is much longer than 30 minutes. So we ended up quantifying the benefits of the automated DAS for those events as well, even though they wouldn't fall within the scope of Branch Technical Position 19.

MEMBER STETKAR: You did this also for all pressurized water reactors?

MR. BLANCHARD: We have a table.

MEMBER STETKAR: We don't have the table for that.

MR. BLANCHARD: Right. We just handed out the table in this presentation for the BWRs. Yes, there's a table for this for the BWRs as well, right?

And the LOCA frequencies themselves come from NUREG/CR-1829, and that is a generic source of data that most PRAs in the U.S. use for their LOCA frequencies.

The large LOCA frequency in there of two times 10 to the minus fifth per year, that's larger than what appears in most PRAs. That is because we expanded the definition of the large LOCA break spectrum in order to define where the automated DAS would be of benefit from the low pressure injection system standpoint.

All right. The number in the upper righthand corner, the common-cause failure of the ESFAS probability.

If we can go to the next slide?

That number we have up there is 10 to the minus fourth per demand. I would like to talk about three things with respect to that number.

First of all, the first thing I would like to talk about is the level of detail in this model with respect to the digital ESFAS system. The second thing I will talk about is the level of detail. I would like to talk about what we considered in the way of failure modes of this particular digital system. Then we would like to talk about where the probability leads you. We will talk about all three things.

What we like to do normally, what I think the vendors of the new plants are doing, and I know that those utilities with current plants that are considering digital upgrades are doing, is model the protection system hardware down to the component level, sensors, communication modules, voting logic. actuation devices, and then assign software failure modes to the hardware that they include in the digital INC model.

For this particular application, we don't

have the details of an INC system, a digital INC system. So we fell back and modeled this at a higher functional level. In this investigation, we modeled the digital ESFAS as a super-component.

You might ask, and we were asked during the meeting a week ago, how could we get any risk insights if we don't model this to detail? What we have done here is modeled this particular digital system at the level of detail we needed to do this application.

We are not trying to make a judgment on this digital ESFAS system in terms of the details of the design, how many channels it has, whether or not it needs watchdog timers for particular failure mechanisms, what its voting logic ought to look like. We are accepting the design of this system and asking the question, what of the effects are of a completely diverse actuation system to this digital ESFAS system?

So we are modeling this system at the level defined by the problem as it is laid out in BTP-19 and ISG 2. So the scope and level of detail of this particular model are commensurate with the definition of the problem for this particular application.

Historically, we have done this for a

number of INC systems. The reactor trip system in most PRAs is modeled at this level. The ATWS rule, when it was developed, had some pretty thorough PRA background developed for the ATWS rule, and the reactor trip system was modeled at the super-component level there. So we are doing something very similar to what we have done in the past.

MR. HECHT: As I recall, you are using a number like 10 to the minus fourth per year?

MR. BLANCHARD: We're going to get there.

MR. HECHT: Okay.

MR. BLANCHARD: Oh, no, no. Per demand.

MR. HECHT: Per demand?

MR. BLANCHARD: Yes.

MR. HECHT: Okay.

MR. BLANCHARD: We'll get there.

MR. HECHT: Okay, fine.

MR. BLANCHARD: Yes, it is in another couple of slides. Let's move to the next slide.

Now what about the failure modes of this digital system? Again, if we had a detailed design, we would like to go to the FMEAs that are generally available for such a safety system. In effect, the failure modes that Bruce presented earlier, and I think the slide number was 24, sensor failures,

communication devices failures, power supply failures, and the failure modes associated with all of those.

Normally, what we would do, if we were building a PRA, we would go back to the FMEA and identify those failure modes. We would also have a D3 evaluation that was developed in accordance with BTP-19 to look at the effects of common-cause failure and get insights out of that, and what to incorporate into our PRA with respect to the common-cause failures.

But given that we don't have these design details, what we elected, once again, was to fall back on the super-component approach, and we made the assumption that, whatever the failure was of this digital system, it failed the components that it controlled in the failure modes necessary for them to fail to perform their function.

Now maybe this digital system doesn't have the failure modes that would necessarily cause the loss of some of those components that it controls and those failure modes, but we simply made the assumption for this application that that, in fact, was the effects of the failure modes of this digital system, and we do have those modeled in the PRA already.

CHAIR APOSTOLAKIS: I don't understand what you just said.

MR. BLANCHARD: Oh, I'm sorry.

CHAIR APOSTOLAKIS: The reason why we have a system there is to open the MOVs.

MR. BLANCHARD: Yes.

CHAIR APOSTOLAKIS: So, if the system fails, it doesn't open the MOVs.

MR. BLANCHARD: That's right.

CHAIR APOSTOLAKIS: So what else do I need to talk about? I mean it doesn't have the failure modes. I don't understand. The failure automatically means it doesn't do its job, which is what an MOV is --

MR. BLANCHARD: A failure to function, the failure modes for this digital system were assumed --

CHAIR APOSTOLAKIS: Yes.

MR. BLANCHARD: -- not to close the breakers --

CHAIR APOSTOLAKIS: Yes.

MR. BLANCHARD: -- and not to open the valves.

MEMBER STETKAR: I think he is saying they did not look at spurious actuations.

CHAIR APOSTOLAKIS: No, that's later.

MR. BLANCHARD: That's later. We did look at spurious actuation.

CHAIR APOSTOLAKIS: That's in the cost calculation.

MR. BLANCHARD: Right.

MEMBER STETKAR: No, no, no, no, no. In mitigating an event, I don't want the system to operate spuriously sometimes.

CHAIR APOSTOLAKIS: And we never do that, actually.

MEMBER STETKAR: Well, that is why we had the problems with the fire analysis, isn't it?

CHAIR APOSTOLAKIS: It seems to me, when we say failure in this case, we mean we don't inject water in low pressure.

MR. BLANCHARD: Ultimately, that's where we ended up.

CHAIR APOSTOLAKIS: Which means these two things.

MR. BLANCHARD: Yes, and that means --

CHAIR APOSTOLAKIS: Well, actually, one of them, either one.

MR. BLANCHARD: Either one.

CHAIR APOSTOLAKIS: Either one.

MR. BLANCHARD: Right.

CHAIR APOSTOLAKIS: So let's go to the meat of it, David, the 10 to the minus 4. Everything

else you are saying, we are on your side.

(Laughter.)

MR. BLANCHARD: Okay.

CHAIR APOSTOLAKIS: Now we are running away.

MR. BLANCHARD: Okay. All right. Now where does the 10 to the minus 4 --

CHAIR APOSTOLAKIS: Well, the IEC you're saying.

MR. BLANCHARD: Yes.

CHAIR APOSTOLAKIS: That is not the Bible.

MR. BLANCHARD: I understand that.

CHAIR APOSTOLAKIS: It doesn't even come close.

MR. BLANCHARD: I understand that.

CHAIR APOSTOLAKIS: So why did you say you were going to do some sensitivities, that is, consider a range of numbers?

MR. BLANCHARD: And that is the last bullet on this slide.

CHAIR APOSTOLAKIS: That's what you did.

MR. BLANCHARD: Yes. So we didn't have the design. We didn't have the vendor operating experience.

CHAIR APOSTOLAKIS: Well, even if you did

have the design, don't tell me you could quantify a software failure.

MR. BLANCHARD: I am being told by vendors of other industries that they can provide me with --

CHAIR APOSTOLAKIS: Those are the guys who publish in the IEEE transactions --

(Laughter.)

MR. BLANCHARD: I understand.

All right, yes, we borrowed 10 to the minus 4 --

CHAIR APOSTOLAKIS: The subject when I started working in INC a number of years back, one of the things they did, they visited places like Boeing, and so on, and the message they came back with was ignore the literature.

MR. BLANCHARD: Okay. Well, we ignored the literature here.

(Laughter.)

With the exception of IEC.

(Laughter.)

CHAIR APOSTOLAKIS: Very good. Very good. Just that exception.

MR. BLANCHARD: Right.

And for the moment, began with a 10 to the minus 4 failure demand probability under the

assumption that this digital ESFAS system would be a high-quality system in accordance with these standards.

MR. HECHT: What you did is you are basically assuming a number and you have done some variation over that number.

MR. BLANCHARD: Yes.

CHAIR APOSTOLAKIS: That's good.

MR. BLANCHARD: And then we will do sensitivity studies on these values to see what impact it has on results.

MR. HECHT: In contrast to actually measuring very low failure rates, it is possible to get some handle on what the probability of failure on demand is in a reasonable time by doing tests and then using I believe it is a Bernouli distribution on the confidence limits.

MR. BLANCHARD: I don't have a system to do that with.

MR. HECHT: I'm just saying it would be possible to determine that.

CHAIR APOSTOLAKIS: Well, we are getting design error. If the cause of the fault is design error, these kinds of things don't really help you.

MR. BLANCHARD: Let's go back --

CHAIR APOSTOLAKIS: Now we have a record here. Whatever you say, and its sensitivity and all that, for your purposes, we may say it is good enough. I don't want other people to come here later and say there is precedent; the ACRS blessed whatever it did in this case; therefore, we are going to do the same. Okay? I hope nobody was going to do that.

MEMBER STETKAR: Could we go stronger and say we don't have any confidence whatsoever in that 10 to the minus 4 number? At least one of us doesn't.

CHAIR APOSTOLAKIS: I second that, too. So at least two of us don't.

MR. BLANCHARD: I understand that.

CHAIR APOSTOLAKIS: But, no, this is very important because people do that. They come back. "But, you know, when Dave Blanchard was presenting, you were so nice to him."

(Laughter.)

Okay. I understand what you are trying to do. You also have a problem with this spurious stuff once in a lifetime. How about Methuselah?

(Laughter.)

So you are doing sensitivity analysis, trying to draw some conclusions, and if that's the best you can do, that's fine. Let's go on.

MR. BLANCHARD: Okay. Can we go back to the table?

MEMBER BROWN: Did you make a stronger statement, like John said, that these are developed --

CHAIR APOSTOLAKIS: His statement is strong.

(Laughter.)

MEMBER BROWN: I'm trying to figure out what he asked, that's all.

CHAIR APOSTOLAKIS: Well, I'm trying to understand.

MEMBER STETKAR: George seconded it.

(Laughter.)

CHAIR APOSTOLAKIS: I seconded it. You can vote, if you like.

MEMBER STETKAR: A straw vote.

MEMBER BROWN: Ten to the minus 4, 10 to the minus 9, you know, they're all numbers. So we all have candy at a child's party.

CHAIR APOSTOLAKIS: That is very true.

Okay, David, what are your conclusions? You are approaching the hour.

MR. BLANCHARD: Multiplying those two numbers together, I have a number on 10 to the minus 9.

CHAIR APOSTOLAKIS: Yes. Yes.

MR. BLANCHARD: Do you want to hear about the spurious actuation?

CHAIR APOSTOLAKIS: Yes.

MR. BLANCHARD: All right. So let's go three slides ahead, four slides ahead.

CHAIR APOSTOLAKIS: Just go to the meat of it. Remember, this is what I've got.

MR. BLANCHARD: This is the spurious actuation frequency and how it was derived.

MEMBER STETKAR: You know, I brought up the spurious actuation a couple of times.

MR. BLANCHARD: Oh, I'm sorry.

MEMBER STETKAR: I'm not so much interested in the spurious actuation frequency per year as spurious action of the safety functions. I am interested in, given a trigger event -- let's call it a LOCA for the moment -- are there spurious actuations of the protection systems that would exacerbate that event, rather than just simply fail the design mitigation functions? Follow me?

That's the problem that we face in the analogy in the fire analysis, that given a trigger event, could things, for example, spuriously open valves? Now if we are talking about a BWR LOCA and it

is a large LOCA --

MR. BLANCHARD: Yes.

MEMBER STETKAR: -- that is not necessarily a bad thing. Okay? But I'm more interested in the transient side of the business and that sense of spurious actuation, not the initiating frequency of spurious safety injection or steam line isolation, or something like that.

MR. BLANCHARD: Okay. I would like to get back to the spurious actuation frequency, but with respect to spurious actuation of the ECCS during another transient in which it wasn't demand, several of us have said we don't normally model that in PRA right now. Part of the reason for that is that we qualitatively truncate it because we don't think it is very likely. Right? The trigger isn't there for it to actuate. It could be, in which we would have a fire, as an example.

MEMBER STETKAR: Right.

MR. BLANCHARD: Yes.

MEMBER STETKAR: Or some as yet undetermined mechanism that would initiate that failure mode.

MR. BLANCHARD: Right. But even that mechanism has a probability of occurring. Thus far in

PRAs, we haven't identified that.

MEMBER STETKAR: We've not examined it. We've not really examined that. That's right.

MR. BLANCHARD: Right. Yes. So, yes, there are other events for which this system could actuate that it would exacerbate. One I could think of in PWR is you might not want to actuate these things spuriously during a main steam line break.

MEMBER STETKAR: Right.

MR. BLANCHARD: Right now, we do, as a part of the ESFAS.

Now with respect to the spurious actuation frequency that was used in this analysis, we interviewed, it says 20 years of LERs. I think that 20 years actually came from Bruce's work. It was more like 15 or 17 years were the LERs here.

The general transient loss of feedwater and loss of main condensor are the categories, and we screened out what I call non-applicable events. What we wanted to do was find every general transient or loss of feedwater, loss of condensor event that occurred as a result of ECCS actuation or actuation of the ESFAS on the secondary side of the plant.

We included in that data collection effort not only just ESFAS-related and initiated events, but

the ATWS system, too. Okay?

We ended up screening some thousand LERs, down to about four dozen. We further screened those out. Why we did that was because of the definition we ended up with on the actuation signals we wanted, the characteristics of the diverse actuation system from an actuation standpoint.

We needed multiple diverse signals in order to actuate this system. That pretty much would eliminate any spurious sensor trips from these 40 or at least two dozen events or four dozen events. So we eliminated those.

Any of these four dozen events that were ESFAS-initiated that was due to a loss of an instrument bus we eliminated.

Maintenance and testing errors, there has been a significant decreasing trend in trips as a result of maintenance and testing. We don't expect the diverse actuation system to be maintained as much as an ESFAS would at power anyway. So we eliminated those.

We ended up with seven events. Those seven events were roughly split between the ECCS spurious actuation events and the events on the secondary side of the plant. Okay?

That works out to be around .005 per year. Half of it is due to the spurious MSIV closures, as an example; the other half is due to spurious ECCS expiration.

Now during the TWG meetings, as we developed this number, we got the comment, "Why don't you just use the historical spurious actuation frequency for the ATWS systems?" In fact, we collected those data, that data. If you go back and look through the LERs we collected for this 17-year period, what you will find is that there are two ATWS-initiated events among these 49 events.

Shortly after this period that we collected the data for, there was another one. So we basically have three spurious ATWS events that caused plant trips over the period of the study here. That happens to be the same number we have generated here as a result of ESFAS. Okay?

We basically have three ECCS actuation events from spurious ESFAS. We have three ECCS -- or excuse me -- MSIV closure events, based on ESFAS. We also have three ATWS-related spurious actuations. Our number wouldn't have changed, had we done the analysis that way. That was the answer that we gave to the Task Work Group.

Now if you go back to the slide one more time, all right, so at the bottom here we have the spurious actuation frequency. We need a conditional core damage frequency for spurious MSIV closures for this reactor trip. We got that conditional core damage probability for each of these plants. That is shown in the upper row of the lower half of the table.

The product of those two numbers gives us the core damage frequency associated with spurious actuation of this diverse actuation system. Then we want to compare that number with the benefits in the top of the table. As you can see, they were roughly on the same order of magnitude. Okay?

So, given the assumptions we have made here regarding probability, the frequencies of the events we are trying to mitigate, spurious actuation frequencies, the actuation or the probability of failure of the digital ESFAS, it is kind of a wash. We are introducing about as much risk as benefit.

MEMBER STETKAR: Dave, can I ask you -- it is less evident for the Boiling Water Reactor. That's why I asked originally about the Pressurized Water Reactor.

You said you looked at the overall risk/benefit from the DAS for a variety of types of

initiating events.

MR. BLANCHARD: Yes.

MEMBER STETKAR: But if I look at the Pressurized Water Reactor analysis, the only benefit you evaluated was for the large LOCA because your conclusion was that DAS, according to the Branch Technical Position 19 assumptions, et cetera, was the one it gave you any protection against.

I don't see a benefit from DAS for Pressurized Water Reactors for transients. In other words, if I run a transient for a Pressurized Water Reactor, and I will come back to my B&W case, where I have really quick steam generator dryout times, if I have no automatic auxiliary feedwater actuation, I should see some benefit from DAS for that type of transient.

Now how large that is, I'm not quite sure, but I don't see that you have evaluated those types of transient benefits.

MR. BLANCHARD: For the PWRs, there already is a diverse actuation system for aux feedwater. We are supporting and endorsing that for ATWS purposes and basically taking credit for it. We are not saying that it is something that shouldn't be done.

MEMBER STETKAR: Oh, okay. So your conclusion for PWRs is that DAS should exist for --

MR. BLANCHARD: DAS does exist and should exist for aux feedwater.

MEMBER STETKAR: Maybe for aux feedwater.

MR. BLANCHARD: It does exist and should exist, yes.

MEMBER STETKAR: I didn't see that on your slide, I guess. I missed it.

MR. BLANCHARD: No, you won't find it on the slides. In fact, I'm not sure it is in the report, but we did not say that you shouldn't automate aux feedwater. Right.

MEMBER STETKAR: It is not as clear to me on BWRs, but you did not look at potential benefits on BWRs from transients either --

MR. BLANCHARD: Yes, we did.

MEMBER STETKAR: -- loss of feedwaters, and things like that.

MR. BLANCHARD: If we can move to the sensitivity study slide --

CHAIR APOSTOLAKIS: I have to go and see a Commissioner. So I will see you gentlemen tomorrow.

MR. BLANCHARD: Okay.

CHAIR APOSTOLAKIS: Mr. Brown will take

over.

MEMBER BROWN: Do I have any special instructions or can I terminate the meeting now?

(Laughter.)

MR. BLANCHARD: We ran a dozen sensitivity studies, and what you are referring to happens to be in one of those sensitivity studies.

MEMBER STETKAR: Oh, okay. Thanks.

MR. BLANCHARD: Under the modeling issues, which is the second bullet on the sensitivity study slide, what you will see for the BWRs is the effect -- let's assume design of the DAS. This is slide 54.

Let's change the design of the assumed diverse actuation system such that it actuates on either of two signals. In the case of the BWR, it would be low level or high containment pressure instead of --

MEMBER STETKAR: There's an "or" rather than an "and".

MR. BLANCHARD: An "or" rather than an "and". What that does now is it brings in the transients as events that could potentially benefit from the automated DAS. Okay? Because during the transients, you can get to a low reactor level all by itself, and that would trigger both the ESFAS and the

DAS.

(Off-mic comment.)

MR. BLANCHARD: That's right. And what happened is that in that sensitivity, yes, we did start addressing some of the transients, but we also raised the risk of spurious actuation.

MR. SIEBER: Yes.

MR. BLANCHARD: And it turned out to be about a wash in terms of the benefits of doing that. Okay?

MEMBER BLEY: I see what you have done on the other one. I haven't looked at this one closely. But, off the top of my head, where we are looking at the benefits, we've got a piece where it seems to me the uncertainty is still pretty high.

MR. BLANCHARD: Yes.

MEMBER BLEY: When it comes to the other one, from the data you used, there's enough data that it is not. On the other hand, a couple of new design DAS we have looked at put the level of effort and protection against spurious actuation well beyond anything I have seen in previous systems.

So, on the one hand, we have uncertainty, either way. On the other hand, it strikes me we are probably pessimistic in how likely we are, at least

for those new designs, to have a spurious actuation. It might still be a wash, but it seems to me there's a --

MR. BLANCHARD: Well, from a spurious actuation standpoint, when we first started doing this analysis, the spurious actuation risk overwhelmed what we are getting in the way of benefits. We kept revising how the DAS actuated and what it actuated until we finally got it down to the point where they were both equal; both the risks and the benefits were equal.

MEMBER BROWN: In the old plants -- I need to ask this question because I got lost in the study as to what. The DAS in the older plants, is that there for ATWS purposes?

MEMBER STETKAR: The only thing that exists in older plants is there is an ATWS mitigation --

MEMBER BROWN: Yes, for a reactor that didn't scram.

MEMBER STETKAR: A reactor didn't scram and feedwater initiation.

MEMBER BROWN: Okay. So for both of those?

MEMBER STETKAR: BWRs, there's feedwater

recirc runback, in some plants feedwater runback, but not all.

MR. BLANCHARD: Right, recirc pump trip and auxiliary rod injection is the ATWS system for BWRs. And PWRs --

MEMBER BROWN: Say that again. ATWS has anticipated transient without scram. I heard all that. I didn't have to deal with that.

MR. BLANCHARD: Right. Trip of recirc pumps, and that lowers power --

MEMBER BROWN: If they trip?

MR. BLANCHARD: No. You trip them deliberately.

MEMBER BROWN: Okay. What's the trigger for telling you you don't have a scram?

MR. BLANCHARD: A high high reactor pressure or a low low reactor level.

MEMBER BROWN: Okay. All right.

MR. BLANCHARD: Either one of those in a BWR.

MEMBER BROWN: Okay. So it is not a power signal. It is another plant signal that tells you something is going on; I shouldn't have seen those, or these are outside the bounds of the transient I would expect. And you use those two parameters to tell you

to trip the reactor?

MR. BLANCHARD: Right. In PWRs, the ATWS system trips the turbine, starts aux feedwater, and in some plants inserts the rods.

MEMBER BROWN: Okay. In BWRs or PWRs?

MR. BLANCHARD: PWRs.

MEMBER BLEY: Pressurized with a drive signal.

MR. BLANCHARD: The signals for that are high high reactor pressure and low steam generator level in most plants. Some plants are different.

MEMBER BROWN: Okay. So slightly different triggering signals, but those are the -- okay.

MR. BLANCHARD: Right.

MEMBER BROWN: This study was, based on the way I read it, was to say, okay, now if we have DAS for other accident mitigation circumstances, the LOCAs, et cetera, independent of this, I don't scram the rods based on these signals other than the BWR gasification where you talked about auxiliary feedwater or something. I'm missing the boat somewhere because I'm not as familiar with BWRs as I am the other stuff.

MR. BLANCHARD: Well, yes, this system we

are talking about here is for actuating the ECCS.

MEMBER STETKAR: I think, to help Charlie out, this is in addition to any separate ATWS mitigation.

MR. BLANCHARD: Right.

MEMBER BROWN: Yes, that's the way I read it, but I heard you guys talking about this other circumstance that it actually takes action on also. That's what I was missing the boat a little bit. I thought you said something was due to some feedwater circumstance. Or do you turn that on? Is it something you do turn on or turn off, or something like that, if it is part of this ATWS mitigation? Or trip? What you said the first time.

I may not be asking the question right.

MEMBER STETKAR: We'll talk later.

MEMBER BROWN: Okay. Let's go on.

MEMBER STETKAR: It is not germane to what we're --

MEMBER BROWN: No, that's fine. I got the second point, and it was fine. I just don't understand the first part of it.

Go ahead.

MR. SIEBER: Well, let me add, in PWRs it's not containment pressure; it's the suppression

pool.

MR. BLANCHARD: I am sorry. The limit on containment?

MR. SIEBER: The containment pressure is the suppression pool --

MR. BLANCHARD: Yes, what happens is you get above 200 degrees; you start approaching saturation, and you effectively have bypassed your suppression pool. Now power is going directly to containment without being condensed.

MR. SIEBER: Okay.

MR. BLANCHARD: That gets to a containment overpressure. That takes about an hour to get there, if you fail to trip.

MEMBER BROWN: I just wanted to make sure I understood that this was to expand the DAS function into the other ECCS-type functional responses to see if it would provide a benefit. I will worry about the other precursors later.

MR. BLANCHARD: Okay. We were talking about some of the sensitivity studies. I guess we discussed that the BWR -- we took a look at modifying the diverse actuation system to actuate on just low reactor level or high containment pressure. That allowed us, then, to take credit for it during

transients, but it also increased the spurious actuation frequency.

For the PWRs, the effects of actuating both high and low pressure systems was examined. Remember, we decided all we needed to actuate was low pressure injection in order to meet BTP-19.

We asked the question, well, couldn't we expand the benefits of this system by actuating high pressure injection systems as well? And the answer to that came out for those BWRs that have a high head, high pressure injection system, what happened is that you increase the challenge to the safety valves if you do that, and their risk goes up significantly as compared to the benefits, if you do that.

MEMBER STETKAR: For the plants that only have the high head injection?

MR. BLANCHARD: Right. Other than that, the risks are about the same, and you get a little bit more benefit out of the PWR.

MEMBER STETKAR: What I missed for the PWRs was the DAS includes automatic feedwater, emergency feedwater actuation on the --

MR. BLANCHARD: Yes, the ATWS DAS has always included that.

MEMBER STETKAR: Okay.

MR. BLANCHARD: It does for this study.

MEMBER STETKAR: Okay.

MR. BLANCHARD: And we are not saying in any way that that's not a good thing to have.

MEMBER STETKAR: Okay.

MR. BLANCHARD: That's right.

MEMBER STETKAR: Okay. Thanks.

MR. BLANCHARD: All right.

The remainder of the sensitivity studies, let me just go over the numerical ones. We did things like set the LOCA frequencies to their upper bound. The conclusions of this study didn't change in terms of the overall benefits being relatively small to begin with.

The ESFAS failure probability, we borrowed the 10 to the minus fourth. We don't really have a distribution for that. So what we did here was simply, rather than set it to its upper 95th percent bound, we said, all right, how high did we have to raise it before we start encountering any of these regulatory thresholds with respect to increases in risk?

We had to raise the failure probability in the actuation system up to around .1 before we started encroaching on the 10 to the minus sixth, 10 to the

minus fifth range, in terms of an increase in risk.

So the conclusion was that, in order for this to have risk/benefit, a good value impact, we have to have a digital ESFAS that we believe is significantly less reliable than what we have in our current analog systems.

We did some completeness examination. We talked about the failure modes of the ECCS. We went back and --

MEMBER BLEY: Tell me again what you -- you raised the failure frequency for the --

MR. BLANCHARD: Failure probability.

MEMBER BLEY: You go after the common-cause failure?

MR. BLANCHARD: That's the common-cause failure probability software --

MEMBER BLEY: Only up to --

MR. BLANCHARD: .1 before we got to the 10 to the minus sixth, 10 to the minus fifth, threshold range that you will find in NUREG/BR-0058 or Reg Guide 1.174.

MEMBER STETKAR: It was less than 10 to the minus 1. It was about eight times 10 to the minus -- quite a bit less. I don't quite understand that because, if I look at your slide, if you go back

to slide 47, if I look at the large LOCA event frequency, it is 2.3 times 10 to the minus 5 per year.

Miraculously enough, the core damage frequency is precisely 10 to the minus four less than that. It is 2.3 times 10 to the minus 9, which says that whatever value you put in there is a direct translator to core damage. So if I keep my LOCA frequency the same at 2.3 times 10 to the minus 5 per year in order to get less than 10 to the minus 6, I have to have something that is about, oh, 4.-something times 10 to the minus 2, not .1. Plus, it's got to be a little bit better than that because it has to mitigate the other LOCAS that come in there.

MR. BLANCHARD: Yes. The actual numbers for the PWR were -- or the ESFAS failure probability had to be .4, and for the BWR it had to be .04.

MEMBER STETKAR: .04?

MR. BLANCHARD: And I summarized in this slide as .9.

MEMBER STETKAR: It has to be 25 times better than this.

MR. BLANCHARD: Yes.

MEMBER STETKAR: I'm sorry. Two and a half times better than this.

MR. BLANCHARD: Right.

MEMBER STETKAR: For BWR.

MR. BLANCHARD: And I think if you go into the sensitivity studies, you will find those numbers, .4 and .04.

MEMBER STETKAR: At that level, what are you approaching for the risk result?

MR. BLANCHARD: For the BWR, we are approaching 10 to the minus 6 per year change in core damage frequency associated with the DAS.

MEMBER STETKAR: That would be the sum of the core damage frequency from large LOCAs and small and medium LOCAs --

MR. BLANCHARD: That's the sum of everything.

MEMBER STETKAR: No, large and small, those LOCAS without anything else? It's only the contribution from those two particular initiating events that you look at?

MR. BLANCHARD: For the small LOCAs, it would be, we were crediting operator action in addition to the DAS because there was significant time available.

MEMBER STETKAR: Yes.

MR. BLANCHARD: Right.

MEMBER STETKAR: But it is only those two

initiating events that you were looking at?

MR. BLANCHARD: Yes.

MEMBER STETKAR: Okay.

MEMBER BLEY: And for the PWR?

MR. BLANCHARD: The PWR approach, the threshold is 10 to the minus fifth, and that would get you at least more failure probability from the ESFAS.

MEMBER STETKAR: Yes.

MEMBER BLEY: There are a bunch of keys to this, but the main key is there's only a couple of initiating events that really make a big difference?

MR. BLANCHARD: Right. They are rare --

MEMBER BLEY: Compared to human action? You looked at all of them, but the real key is, for most accident sequences in the PRA, manual action is plenty good enough?

MR. BLANCHARD: Plenty good enough. Exactly right.

MEMBER BLEY: That's the real --

MR. BLANCHARD: And there's only a handful for --

MEMBER BLEY: For common-cause failure?

MR. BLANCHARD: Right, and there's only a handful for which the automated DAS is of potential benefit, and even those are relatively small.

MEMBER BLEY: Now I understood there was a manual --

MEMBER STETKAR: I don't think they're saying you don't want a manual stop.

MR. BLANCHARD: All right.

MEMBER BLEY: I'll say it out loud. It leaves manual DAS is the way to get around failures from the standpoint of, if we are uncertain and uncomfortable with that as a defense in-depth --

MR. BLANCHARD: From a defense in-depth and diversity standpoint, we are actually doing more than just manual DAS. That is what this next slide, slide 52, is going to be about.

Why did the numbers come out the way they are? It is because there's effective defense in-depth and diversity provided already by existing plant features. That defense in-depth and diversity is first provided by adequate protection against the occurrence of a LOCA through designing the primary coolant system in accordance with piping and pressure vessel codes, periodically inspecting the primary coolant system in accordance with Section 11, performing hydros as we start up before we ever go to power following outages, for fueling outages. Then monitoring the performance of the primary coolant

system with leakage detection systems during power operation.

By these leakage detection systems I am referring to sump-level monitoring, sump pump operating times, new cells in the containment, radiological concentrations in the containment atmosphere, those types of leakage detection systems.

MEMBER STETKAR: And those are part of the reason why we have the small frequencies that are used in the PRA.

MR. BLANCHARD: Exactly right.

We also have a highly-reliable ESFAS. That is due to all the design in accordance with existing standards, the fact that there is redundancy and independent trains associated with that ESFAS. For software, rigorous validation/verification programs, and perhaps even design features such as the defensive measures that we discussed earlier today that limit the potential for those INC failures and common-cause failures.

And finally, these two defense in-depth measures, if you will, are independent of one another. The introduction of a software failure can cause the LOCA. A pipe flaw is not going to cause the introduction in the software error into the digital

ESFAS. There are independent root causes that cause each of these failures.

MEMBER STETKAR: What about spurious ADS on a boiler?

MR. BLANCHARD: I haven't had anybody ask me that question for a year.

(Laughter.)

There is an answer to that.

MEMBER STETKAR: Okay. Maybe if you come back tomorrow --

MR. BLANCHARD: Okay, yes, let me think about that. That question has come up, and it has come up especially with respect to the passive plants.

MEMBER STETKAR: Yes. I was thinking ahead to the passive plants.

MR. BLANCHARD: And it is with respect to the AP1000 in addition to the ESBWR. So, yes, I'll think about that again.

MEMBER STETKAR: Think about it. I'm interested in the answer to that one.

MR. BLANCHARD: Okay.

MEMBER STETKAR: Because you have done the studies out of the existing -- think within the context of the existing fleet.

MR. BLANCHARD: Right.

MEMBER STETKAR: It is sort of an interesting question, at least to me.

MR. BLANCHARD: Right.

Okay. We also have this conclusion that, with respect to spurious actuation, we may be introducing about as much risk as we are taking care of here. What's the reason for that?

Well, first of all, this diverse actuation system is intended to mitigate events that are fairly rare, large and medium LOCAs, and aren't expected to occur in any plant over the life of the entire fleet. And given what we came up with as a spurious actuation frequency, we might be causing shutdown of a plant every several years. We may be tripping a plant every several years to address an accident that may not occur over the life of the entire fleet.

Now these inadvertent shutdowns are design-basis events. They are covered in the design basis, but they are not without risk. Okay?

And they occur at a significantly higher frequency than the events we are trying to mitigate. So that's the reason why we end up introducing perhaps as much risk or even more than we might be mitigating here.

The outcome of that is, if you decide to

put in the diverse actuation system, you really need to take a look at making sure it's robust against spurious operation, perhaps even more robust than the existing ATWS systems.

All right. And in that regard, one of the outcomes of the analysis is, should you decide to put in the diverse actuation system, here are some of the characteristics of that diverse actuation system that fell out of our analysis. We have talked about some of them because they were assumptions we made going into the accident sequence quantification.

We want their actuation to be based on multiple plant conditions and we want to have all those conditions before we actuate it, both pressurizer pressure and high containment pressure in the PWR, low reactor level, and high drywall pressure in the BWR.

We want to require power to actuate this system. We don't want a loss of an instrument AC bus to actuate the system. That is similar to what we have right now for the ATWS systems.

We don't want any LCOs or allowed outage times to cause a shutdown of the plant, the manual shutdown of a plant. If you shut down the plant once during the life of the plant, the benefits are so

small, you have pretty much introduced enough risk to wash out any of the benefits.

Rather, we are proposing putting availability and reliability requirements and monitoring the performance of this system along with all the other systems that are in the Maintenance Rule. The ATWS are similar in that regard as well.

Next slide.

We don't believe we need an automated DAS for the steam line breaks. This says downstream, the MSIVs. That's for the BWRs, the steam line breaks inside the containment for the PWRs. That's because we have significant time available for the BWRs, and the reactor coolant system and fuel conditions appear to be more benign in the PWRs if we don't actuate the ESFAS.

Now everything up to this point was considered in the existing analysis. There were a few additional design characteristics that came out of the TWG activities that we didn't have time to evaluate, but we have put them here in the list anyway for future consideration.

One way to limit spurious actuation in the system is, in addition to the characteristics that we have already defined for the system, is put in series

with the diverse actuation system a timer set at the latest time necessary to initiate the system and still have it be effective in providing adequate core cooling, based on best estimate from a hydraulic analysis.

In other words, if you are going to actuate low pressure injection with this system, put a two-minute timer on it. Give the operator a couple of minutes' chance.

MEMBER BLEY: And a bypass switch for the operator?

MR. BLANCHARD: And a bypass switch, just like the ADS timers that we currently have in the BWRs. In the event it should spuriously actuate, a few minutes may be enough for him to be able to recognize now this isn't anything that needs to be going off at this point.

MEMBER BROWN: So if the timer fails, you block it?

MR. BLANCHARD: No. If you have an ADS or if you have a DAS actuation, the timer starts, and it gives you a few minutes to assess whether or not --

MEMBER BROWN: Before it actuates?

MR. BLANCHARD: Before it actuates. And if, in fact, you don't have these conditions --

MEMBER BROWN: If you have a timer failure, it doesn't actuate then? You've introduced another failure mode into the --

MR. BLANCHARD: You could have entered a failure mode.

MEMBER BROWN: That ends up with no diverse actuation system operation?

MR. SIEBER: That is no worse than not having the system.

MR. BLANCHARD: You put redundancy in there, such that you would have --

MEMBER BROWN: You add some more stuff?

(Laughter.)

MR. BLANCHARD: Yes, yes.

MEMBER BROWN: I'm getting down to the bottom line here. You add some more stuff.

MR. BLANCHARD: We could add some more stuff.

MEMBER BROWN: You could just keep adding stuff to make sure we have other modes that we can evaluate to determine whether their failures --

(Laughter.)

MR. BLANCHARD: For the high pressure injection system, you could set this timer at 15 minutes or 20 minutes, whatever is your best estimate

analysis.

The other thing that came out of the TWG discussions was one of the reasons that ESFAS, spurious ESFAS, now causes reactor trips is because it does more than just start the ECCS. For many plants, it trips or isolates systems that you need to keep the plant running, and you load-shed non-safety-related buses, which contain balance-of-plant systems that are required for operation. You isolate non-critical service for our headers to the balance-of-plant. Right?

If you could do a best estimate analysis of the systems that you want to actuate with this diverse actuation system, without those isolation features, without those load-shedding features, and convince yourself that, well, I can actuate these trains of ECCS without all of the other stuff that I normally actuate with ESFAS, in isolating non-critical service water or load-shedding, non-safety buses, then we can overdo a spurious actuation frequency.

We didn't have time to evaluate the benefits of that as a part of this evaluation, but it is something worth considering.

MEMBER BROWN: So the ESFAS is what does the isolation of the other systems?

MR. BLANCHARD: Yes. And you would have to do with that the diverse actuation system, for example, if your ECCS pumps required service water cooling, and you didn't have enough of that unless you isolated the non-critical service water header, you would have to have the diverse actuation system isolate the non-critical service water as well.

If you needed to trip the drywall coolers in a BWR because you have to do that in order to bring the ECCS pumps on and keep the voltage in the plant at levels that the pumps will operate appropriately, then you would have to do that with the diverse actuation system, too. But maybe you could do a best estimate evaluation that said you could run the ECCS pumps without load-shedding the drywell coolers.

MEMBER BROWN: If the ESFAS system actuated without being triggered by its triggering signals, in other words, it's just a spurious actuation of that system over five minutes' time, we'll say, on its own, does that also end up doing all those other things or does it only do that if it is triggered by the --

MR. BLANCHARD: If it is triggered by the ESFAS, yes, it does. If a pump just starts, the low pressure injections won't inject to the reactor

because the reactor pressure is too high. The high pressure systems can.

MEMBER BROWN: I'm talking about there's certain signals that trigger the ECCS ESFAS systems to actuate.

MR. BLANCHARD: Right.

MEMBER BROWN: If they aren't present and you had a spurious start of the high pressure/low pressure injection, whatever they are, the pumps, whichever, then all of these other balance-of-plant stuff, they stay online?

MR. BLANCHARD: They stay online, right.

MEMBER BROWN: Okay.

MR. BLANCHARD: However, the high pressure injection systems can inject to the reactor --

MEMBER BROWN: I understand that, the low pressure --

MR. BLANCHARD: And now you can have reactivity events. MEMBER BROWN: Yes, I understand that. I am just trying to get -- which I will forget tomorrow, but I just wanted to have a calibration today.

MR. BLANCHARD: All right.

MEMBER BROWN: In a few months, I might remember this.

MR. TOROK: Okay, I think we can close this out with a real quick review of the conclusions and recommendations. I'm not going to try to repeat everything Dave said there, that's for sure.

But the bottom line here is we think, based on all that Dave is showing you, that right now with current PRA techniques, it is possible to generate useful risk insights.

Now, in regard to the particular analysis he did, the bottom line was that, for the events analyzed, that the automated DAS shows little to no benefit. I am not going to reiterate the reasons why. Dave just did that.

Another conclusion, based on all this analysis, is that, in general, the high-frequency events are going to benefit more from an augmented DAS than rare events.

The things you just talked about, one of the factors here was the events that the DAS addresses are rare events, right, and it drives the benefit down?

So that is a very brief summary of the conclusions from the report.

Then the resulting recommendations are we would hope that the ACRS will consider what we have

done here and encourage both staff and industry to take advantage of PRA methods where it makes sense. We talked a lot about how you know when it makes sense in terms of sensitivity to assumptions, and so on.

Also, something coming out of the evaluation is that it may make sense to take another look at the BTP-19 and, effectively, let it consider both frequency and consequences in assessing defense in-depth and diversity in the plants, which really means allow a graded approach where the solutions and the protective measures are proportional to the risk. So take advantage of the risk insights is what we are saying here and consider both frequency and consequences.

Finally, I guess, we would like to ask for your concurrence, and we should promote methods for addressing digital system issues that credit both prevention and mitigation. We shouldn't be talking just about what happens after the CCF or what happens after the accident. We should be talking about preventing the accident, crediting both prevention and mitigation techniques.

Where this leads us is back to this kind of leading statement that we had at the beginning of the talk, the DAS talk, which is that there are a

number of ACRS statements that have been made in literature, and construed to mean, more or less, that what we have done here isn't appropriate or isn't possible, or that sort of thing. So we would appreciate clarification of those statements, in light of what we have presented here.

That's about it. Now, beyond that --

MR. SIEBER: You may want to repeat that tomorrow.

MR. TOROK: Okay.

MEMBER BROWN: Yes.

MR. TOROK: Very well. I have no problem. I would be happy to.

This is just the same slide we had at the beginning. These are the major points in regard to operating experience, failures modes, and PRA insights. At this point, I don't see any point in going over them again, unless somebody here wants to.

Again, we have repeated here at the end the same thing we said at the beginning in terms of OE. It would be nice to keep gathering data and getting generally some common definitions. We think it is important to credit defensive measures in regard to protecting against CCF, and we think it is a good idea to use risk insights when we can.

This is kind of a motherhood thing, right? We would like to encourage increased technical exchanges here with us and NRC Research, and so on. I think there have been a number of examples here where it is clear, at least it is clear to me, that it would have been beneficial to have that interaction early on, especially where we were looking at the OE and how we found those events, how we evaluated them, what terms and definitions we were using, why we were using them. We were never really able to have those discussions. It would have been helpful to all of us, I think, if we had. So we would like to encourage more of that.

MEMBER BROWN: What prohibits discussions?

MR. TOROK: Well, apparently, you're going to have to ask --

MEMBER BROWN: Are you trying to stick a stick in the staff or --

MR. TOROK: No, no, no. All we are saying is that we think more of that technical exchange on these issues is a good idea, and we just think we should keep doing it.

We effectively, while we were generating the results from our OE study, we tried to get staff engaged to discuss things that we were doing and why,

and so on, but I guess they have more restrictive rules on engagement than we do. I think that prevented an exchange that would have been really valuable, technical exchange early on.

MEMBER STETKAR: I think we will probably hear more about that tomorrow.

MR. TOROK: We have been working on that. In fact, that's exactly what Dan is pushing for, and I think we are finally making some progress now.

MEMBER BROWN: It's just got to be a coordinated thing. I mean I understand the problem is that they are a regulator and you all are an industry representative evaluating things, and they can't be seen as being in bed with you. Therefore, they want you to develop conclusions, based on your evaluations.

There's nothing that stops you all from obtaining and capturing more of OE information. You don't need permission to go do that, if I'm not mistaken. Correct?

MR. TOROK: We weren't looking for permission. We were just looking for opportunities to discuss the rationale for how we were doing our evaluations. Because we thought it was important for them to see that, and we knew that they had been looking at events and thought they might have some

valuable input along those lines.

Now when I look at the comments we have received back from them, I think some of them are direct results of not having that communication early on because they reflect misunderstandings of what we were doing and why we were doing it. That's all.

MEMBER BROWN: I don't know. It's a tough road to walk between the regulator and the industry.

MR. TOROK: Yes.

MEMBER BROWN: When you engage and when you don't, and under what circumstances. Okay? So I understand.

MEMBER BLEY: You now have the MOU. The key is to get that done and get it agreed with, so that you know what the boundary conditions are.

MR. SIEBER: I guess overall I have trouble with the operating experience because it is harsh and really applies to just certain systems in the United States. I think that needs to expand.

On the other hand, I appreciate the risk analysis work on the diverse actuations. I think that was well done. Thank you.

MEMBER BLEY: I would like to express my appreciation for the whole day. I think we have learned a lot, and you have given us a lot to think

about. A lot of the things you presented are very helpful. To me, rather than being conclusions, they are a real good start for addressing some of the issues. I really appreciate it. I've got to think more about the last stuff you presented. It's quite interesting.

MEMBER BROWN: I am sure the PRA stuff will --

MEMBER BLEY: There might be something hanging around there that I didn't get in the first couple of passes.

MEMBER BROWN: John?

MEMBER STETKAR: Yes, I think I would echo Dennis. I think you have done a tremendous service. I think it provides a framework for thinking about the problem in terms of the last stuff that we saw.

The only caution in terms of the operating experience, the thing that struck me is that at times we can suffer from too much emphasis on classifying events. My only concern about that is that users of the database should not be discouraged from looking at particular events simply because they have been put into a certain classification.

In other words, I wouldn't like to see the classification process steering people away from

certain types of events and only focusing them on the ones that the particular analysts or the particular classifiers felt were most applicable for their purposes.

So, the compilation of the operating experience and the uniform kind of attempts to describe the events is a wonderful resource. It should provide a common basis and a common library for everybody to use in terms of understanding the experience. It is just be a bit careful to not overdo the classification process because it could backfire.

MEMBER BROWN: Myron?

MR. HECHT: I have two comments with respect to the operating experience evaluations. That is, No. 1, how are the data collected? I have seen situations, not in this context but in other contexts, where you can have a set of experience, but not having all the failures, relevant failures, that are being collected can lead to problems. So we have the issue of completeness.

And the other one is during the break I think when we had discussions with Thuy and with you, Charlie, one of the things that I realized is that, in addition to knowing how things fail, we have to know what the things are, which leads to the question of,

what are the designs and the architectures of the systems that we are dealing with?

That leads to another question, which is, in addition to knowing about the failures and classifying the failures, we may need to have a way of representing the systems, so that we know what it is that we are talking about. We can understand what it relates to what.

I would just point out that there are several architectural representation languages and design representation languages for software. There's UML, and at the system level there's SysML and also my personal favorite, which is AADL, which is the Architecture Analysis and Design Language.

But, in any case, as we speak about operating experience, particularly for the digital systems, we kind of know what the plants are. We all have a fairly good idea of what the plants are. That was very clear in the discussions as we were talking about the various kinds of relief valves and the various pressure levels. We all have, I think, a common conceptual picture about that. I don't think we have the same common conceptual picture about the computer systems that are used on the control side and in the safety side.

MEMBER BROWN: I actually thought the discussion on the OE side and the discussions on CCFs were extremely valuable.

Did you have something else you wanted to say?

MR. SIEBER: No.

MEMBER BROWN: Okay. He was pointing, and I didn't want leave him.

I thought it was one of the better discussions of CCF with examples and the ability to discuss architectures and the context of the common-cause failures you all thought about. I thought that came out pretty well. I thought it was a very well-rounded set of discussions on that. I hadn't heard that in some previous meetings.

We were able to wrap in some architecture and some fundamentals in terms of philosophy into that relative to the independencies and the dependencies or non-dependencies, or whatever you want to call them in terms of the architectures. So I thought that was very good.

I am not going to comment on the PRA stuff. I will let the PRA experts deal with that.

The one thing I wanted to springboard from John's comment on the classification because you did

classify no CCF, no software, whatever it was. There were some inherent classifications there.

Some of the events were very serious. So there is a severity issue that you have to look at. If you lose sight of the severity and you say, well, gee, did we really classify it right, I don't know how to mix and match those. I just think you've got to track which ones really were significant.

Now we have made a judgment as to how it is classified, but there may be some other circumstance in which the severity may make you look at it from a different standpoint.

So, other than that, I wanted to thank you all for the presentation today. We will have wrapup from the staff tomorrow morning, give them their day in court. Whether that results tomorrow morning, we will see how that plays.

Other than that, are there any other comments?

Jack?

MR. SIEBER: Yes. I would like to repeat your comments about previous ACRS letters and how new information relates to those letters; you should repeat those tomorrow.

MEMBER BROWN: Oh, yes. Good point.

Thank you. I will make sure we do that. That's great.

I guess I would like to thank you guys for a couple of things: all the time you gave us certainly. I know that is a very big thing for this body to make so much time available for this kind of discussion.

And the other thing, I have to tell you, as an EPRI guy, we often wonder if anybody is reading our stuff, and you guys I thought showed an excellent knowledge. There's evidence that you were looking at it in great detail and you had a lot of really good comments and suggestions as a result of it. We don't get that every day. So I would really like to thank you for that.

MEMBER BROWN: Go ahead, Dennis.

MR. SIEBER: You can enjoy your high tonight.

(Laughter.)

MEMBER BLEY: I want to just mention, and I suspect you have run into it back home, I know a lot of traditional nuclear engineers, safety engineers, who don't do PRA-related things who, when you start having that discussion you went through today, go ballistic, saying, "There's no risk from tripping a

reactor ever, and that's nuts to even bring that up. We are worried about a situation that could be very bad where we need to trip it and don't tell me there's any risk over on that side."

I think you have had a favorable audience from that point because all of us have been thinking about it, but there's not an insubstantial population of nuclear engineers who would challenge that a lot, and we might see something one of these days here.

MEMBER BROWN: Because actuations can cause problems. That is what it sounded like.

MEMBER BLEY: If it a spurious actuation in the system, they wouldn't believe it. They would believe it a little bit for ESFAS and SI, but not for reactor trip.

MR. SIEBER: It provides opportunities for additional adventures.

MEMBER BLEY: It certainly does, but they would say, how can it be bad? It takes a long, long time to convince them that there might be something there.

MEMBER BROWN: Okay. If there's no other comments, we will adjourn until tomorrow morning.

(Whereupon, at 5:17 p.m., the proceedings were adjourned for the day, to reconvene the following

day, Thursday, August 20, 2009.)

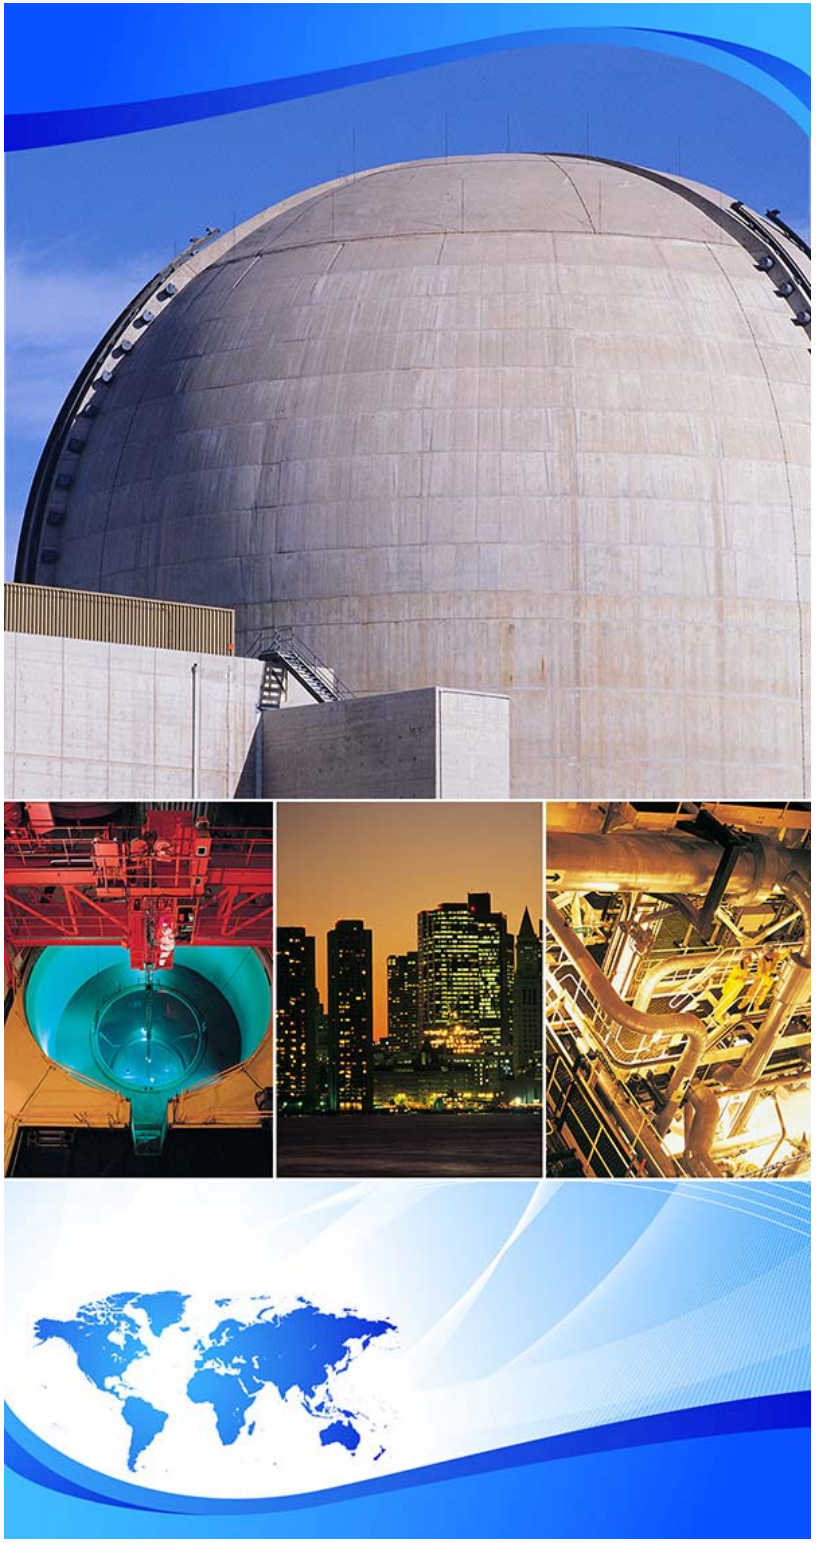

ELECTRIC POWER RESEARCH INSTITUTE

EPRI Projects on Digital I&C

- Operating Experience Review
- Failure Modes
- PRA Insights

ACRS Subcommittee on Digital Instrumentation & Control Systems
August 19, 2009

Rob Austin, Ray Torok
EPRI

Bruce Geddes
Southern Engineering Services

N. Thuy
EDF R&D

Dave Blanchard
Applied Reliability Engineering

Purpose / Topics

Discuss EPRI digital I&C activities

- **Operating Experience Review**
  - *Operating Experience Insights on Common-Cause Failures in Digital Instrumentation and Control Systems* (EPRI 1016731, Dec 2008)
- **Digital Failures - Mechanisms, Modes and Effects**
  - *Common-Cause Failure Applicability* (white paper prepared for NEI Digital I&C and Human Factors Working Group, Feb 2008)
- **PRA Insights**
  - *Benefits and Risks Associated with Expanding Automated Diverse Actuation System Functions* (EPRI 1016721, Dec 2008)

Provide information requested by ACRS in March/April 2008 meetings and May 2007 letter

Gather feedback that will help guide future EPRI research

Purpose of EPRI Research on Digital I&C Issues

- **Provide the technical bases and guidance** to help utilities:
  - Manage I&C obsolescence
  - Implement advanced I&C and information technologies in nuclear plants
  - Enable plants to use digital technology capabilities to:
    - Maintain safe operation
    - Enhance reliability
    - Reduce operating costs
  - **Address regulatory issues regarding digital systems**

Next Steps on EPRI Digital I&C Activities

- Document existing PRA scoping and sensitivity studies
- Publish guidance on protecting against CCF
- Develop guideline on estimating digital system reliability based on design and process attributes (defensive measures)
- Develop guideline for failure analysis of digital systems
- Continue support of NEI Working Group
- Continue activities under MOU between EPRI and NRC Research on digital I&C issues, e.g.,
  - Operating experience
  - Risk methods
  - Adequate diversity and defensive measures for CCF protection
  - Human factors

Context of Research – Support NEI Working Group on Common-Cause Failure (CCF), Defense-in-Depth and Diversity (D3)

Current NRC guidance on CCF / D3 – policy and positions

- **Policy - SECY-93-087** and **SRM**
- "What-to-do" guidance to comply with policy - **BTP-19**
  - Branch Technical Position 19 of Standard Review Plan Ch 7 – *Guidance for Evaluation of Diversity and Defense-in-Depth in Digital Computer-Based Instrumentation and Control Systems*
- Detailed guidance and technical basis - **NUREG/CR-6303**
  - *Method for Performing Diversity and Defense-in-Depth Analyses of Reactor Protection Systems*, 1994
- Clarifications, 'HOV lane,' '30 minute criterion' - **ISG 2**
  - *Task Working Group #2: Diversity and Defense-in-Depth Issues, Interim Staff Guidance*, 2007-2009 (30 min criterion modified in 2009)
- Staff positions - **SECY-09-0061**
  - Includes comments on EPRI white papers

Context, cont'd
CCF Guidance for Postulated Accidents and Anticipated Operational Occurrences (AOOs)

Applicability of EPRI research in CCF evaluations

- Policy - Identify CCF vulnerabilities, ensure adequate diversity (deterministic, prescriptive)

Insights from OE and failure modes research helpful here

- Demonstrate compliance with acceptance criteria of BTP-19
  - Demonstrate adequate diversity, OR
  - Identify corrective actions, OR
  - Provide basis for taking no action

Risk insights applied here

Evaluation approach & results conform with regulatory policy

Key Points

Operating Experience (OE)

- Software has been no more problematic than other CCF contributors
- Need to capture and promote **process** and **design** characteristics that have been effective in protecting against CCFs

Understanding "Digital" failure modes

- "Failure mechanisms produce failure modes which, in turn, have certain effects on system operation" (i.e., failure modes are understandable)
- PRA models represent failure modes/effects, and do not need exhaustive treatment of low level digital failure mechanisms to generate useful insights
- Failure mechanism prevention and mitigation remain very important in designing robust systems (fault avoidance and fault tolerance)

PRA insights

- Risk insights are possible today using existing techniques
- Should encourage use of PRA given its capabilities and current state of the art

Request ACRS Concurrence

Staff and Industry should:

- Continue to gather and apply OE lessons on failure causes, corrective actions and preventive measures – develop common definitions for binning and evaluating events
- Develop methods for crediting defensive measures in protecting against failures and CCF (especially where they are better than diversity), and in assessing digital system reliability
- Use current risk methods to address digital I&C issues for both operating and licensing applications where appropriate, e.g., for low frequency events
- Increase technical exchanges to resolve issues more effectively and efficiently (particularly with RES)

First Topic:
Operating Experience Review

- White paper version presented to ACRS in March/April 2008
  - 322 safety and non-safety events in U.S. over 20 years
  - Look for actual and potential common-cause failures (CCF)
  - Success stories not included – did not look at overall impact of digital
  - Capture insights on causes, corrective actions, defensive measures
- Final EPRI report (1016731) published December 2008
  - Provided to ACRS and NRC January 2009
  - Expanded discussion of methods and observations
  - Appendix with brief descriptions of all 322 events
  - Detailed review by EPRI, NEI Working Group and various technical experts
- **Today will:**
  - **Recap key points and conclusions**
  - **Expand upon discussion of digital failure modes in OE**

Key Terms

- **Defect** – A deficiency in characteristic, documentation or procedure. In software often referred to as "fault" or "bug."
- **Common defect** –
  - Safety Systems - A defect that affects multiple redundancies, for example a software fault that exists in all divisions of a redundant safety system.
  - Non-safety systems – Also includes defects in shared resources, for example a power supply that feeds multiple non-safety process controllers.
- **Trigger** – A plant condition or specific set of inputs that activate a defect; in digital systems this is typically an unanticipated, unexpected, or untested condition.
- **Concurrent triggers -** Triggers which occur over a time interval sufficiently short that it is not plausible that resulting failures (due to a common defect) would be corrected
- **Failure** – Degraded or terminated ability of a functional unit to perform a required function. A software failure results when a software defect is activated by certain triggering conditions.
- **Potential CCF** – A defect common to multiple redundancies that can result in an actual CCF in the presence of concurrent triggers.
- **Actual CCF** – A malfunction on demand that results in an incorrect response or loss of function across multiple redundancies at the same time due to a common defect.
- **Digital event** – Any plant occurrence that involved or affected a digital system and was reported in the databases that were searched.
- **Software event** – An event involving design defects introduced in the software development process (not, for example, incorrect setpoints or flawed requirements)

Key Terms Comparison Chart

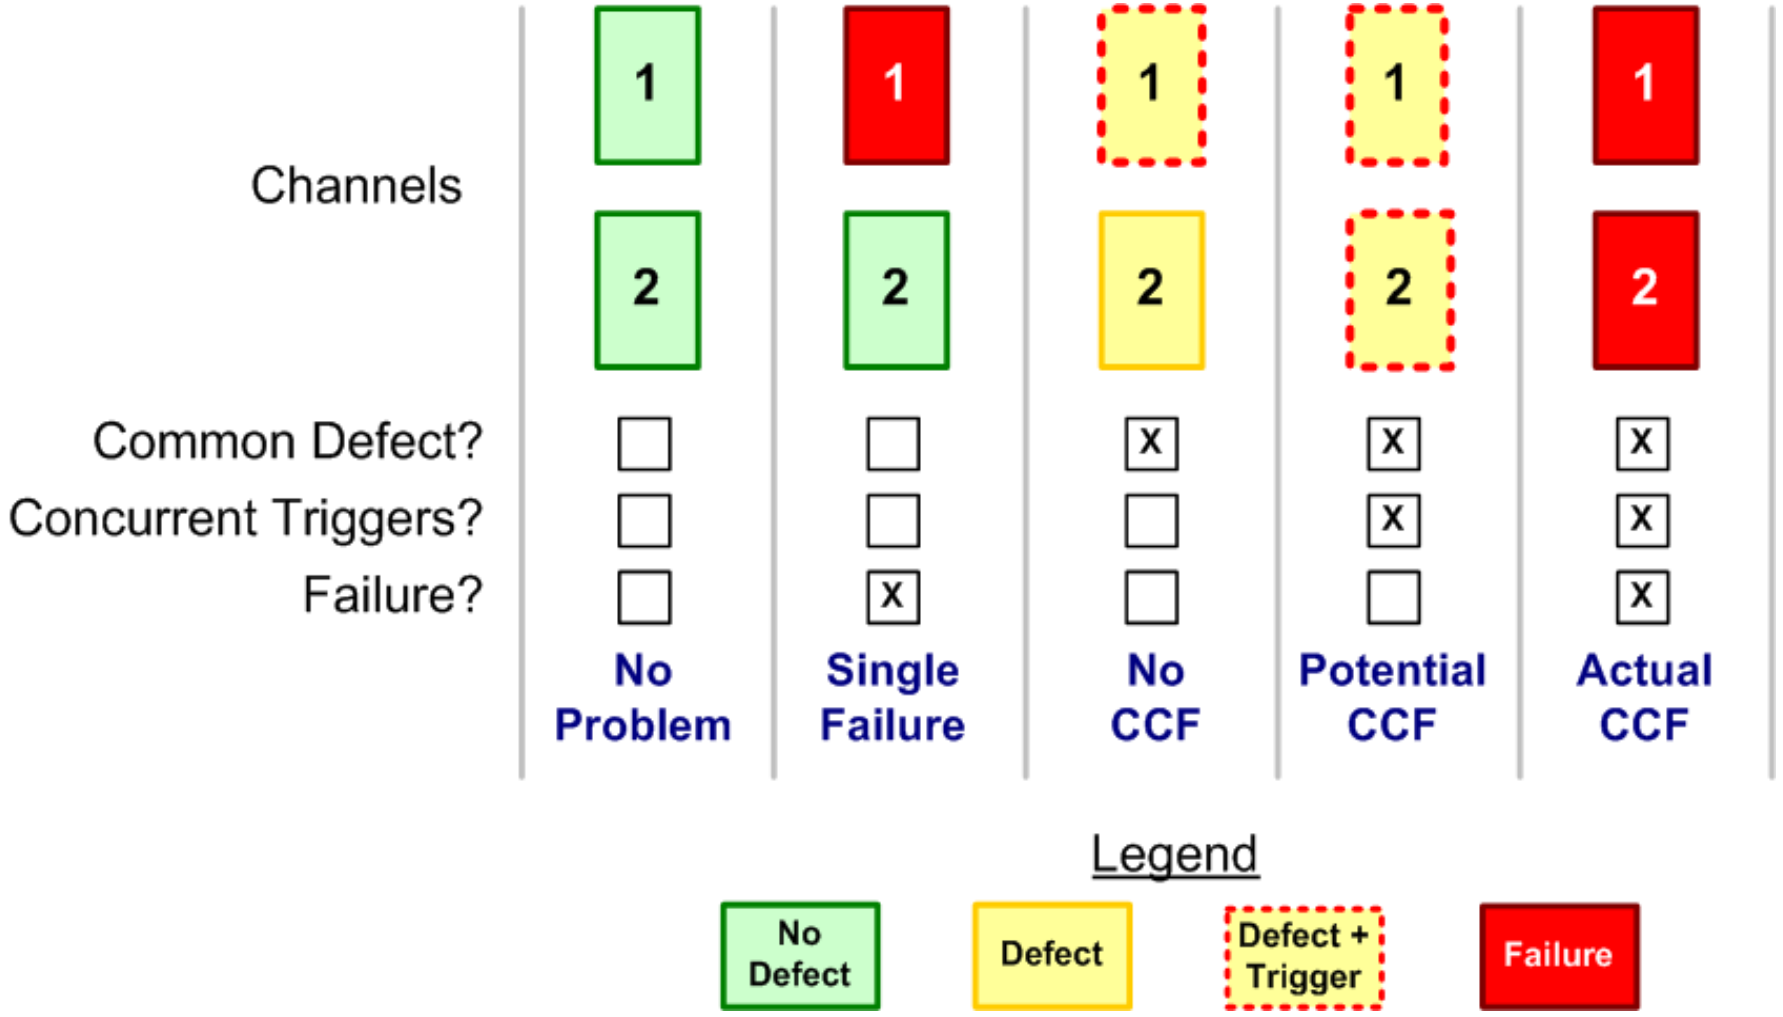

What is the OE Telling Us?

- **There were no actual CCF events that disabled a safety function**
- **Actual and potential CCF events were dominated by non-software issues, e.g.,**
  - Lifecycle management and human performance errors (e.g., incorrect setpoints)
  - Hardware failures (non-1E)
- **OE suggests that current methods are effective in keeping software a minor contributor to CCF**
  - Use of software codes and standards
  - Design and process characteristics that preclude or limit CCFs ("defensive measures" and diversity attributes)

1E Common Defect Events

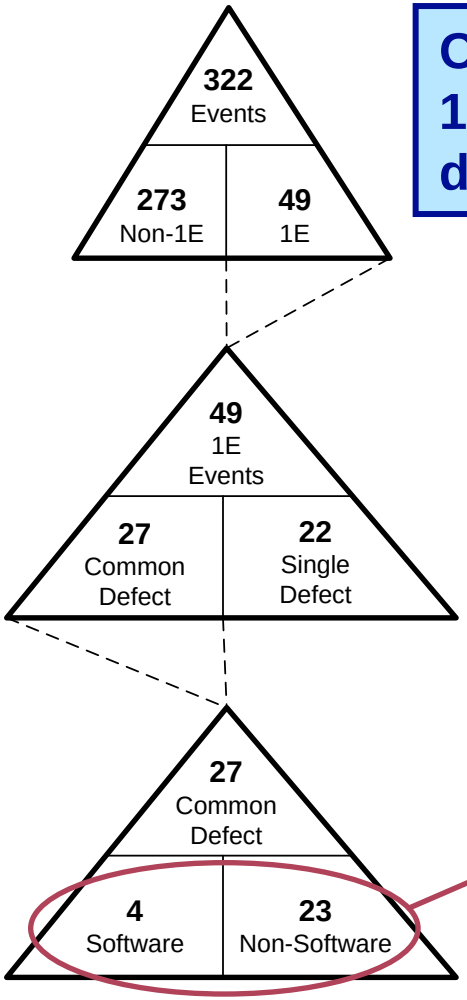

Out of 27 common defect events, 1 could have resulted in a CCF due to software (3.7%)

		1	1
		2	2
Common Defect?		X	X
Concurrent Triggers?		X	X
Failure?			X
		Potential CCF	Actual CCF
System Level	Software	1	0
	Other	5	0
		6	0

Remaining Common Defects:

- 10 Single Failures (**2 due to SW**)
- 6 Spurious Actuations (**1 due to SW**)
- 4 Subsystem Potential CCFs (0 due to SW)
- 1 Subsystem Actual CCF (Setpoint Issue)

EPRI ELECTRIC POWER RESEARCH INSTITUTE

Failure Mechanisms, Modes and Effects in 1E Software Events

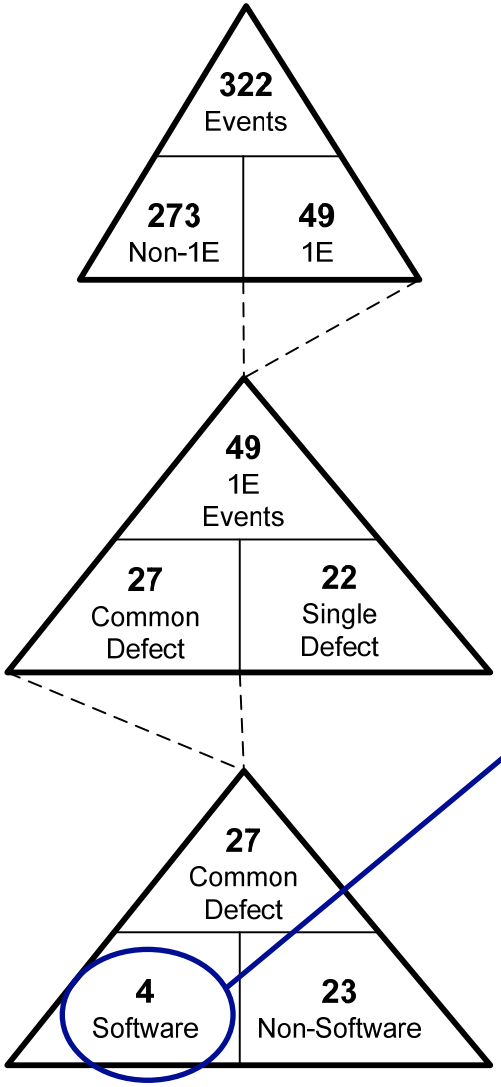

Event	Root Cause	Failure Mechanism[1]	Failure Mode	System Level Effect
1	Specification Error	Incorrect Substitute Value for Failed Sensor **(Task Incorrect Response)**	CPC[2] Channel May Not Trip When Required	**No CCF**
10	Design Error	Incorrect Logic While in Self-Test Mode **(Task Incorrect Response)**	Sequencer Blocks Safety Injection ~ 15% of Time	**Potential CCF**
13	Missing Requirement (Omission)	No Watchdog Timer (HW) & "WRITE" Operation (SW) **(Task No Response)**	RMS[3] Processor Lockup During Power Transient	**No CCF**
221	Design Error	Counter Not Initialized at the Right Time **(Task Incorrect Response)**	Momentary Step Change in RMS Output Signal	**Spurious Actuation**

1. As described in ACRS Letter dated 4/29/08
2. CPC = Core Protection Calculator
3. RMS = Radiation Monitoring System

Event 10

Event #	10	Event Date:	Nov-94	System:	ESFAS

Inoperable Load Sequencer

	System	Subsystem	Channel		
Single Failure				Root Cause:	Inadequate Software Design
Spurious Actuation				Contributing Cause:	Inadequate Software V&V
Potential CCF	X			Contributing Cause:	---
Actual CCF				Corrective Action 1:	Software Change
				Corrective Action 2:	Software Development Process Change

Failure Mode:	Software logic defect in the application code on asynchronous channels can prevent valid safety injection signal from passing through some of the time when in automatic test mode.
Risk Significance:	Auto SI function available 90% of time. Manual actuation available as a backup (SGTR, Small & Med LOCA). Simulator verified manual action could take place in time for Large LOCA

RISK COLOR

Event 10 (Risk Significance)

Initiating Event Frequency	Mechanical System Designs	≥ 3 diverse trains OR 2 redundant systems	1 train + 1 system with redundancy	2 diverse trains	1 train + recovery of failed train	1 train	Recovery of failed train	None
1 to 10^{-1} / yr	Reactor trip Loss of Condenser							
10^{-1} to 10^{-2} / yr	Loss of off-site power Total loss of main FW Stuck open SRV (BWR) MSLB (outside cntmt) Loss of 1 SR AC bus Loss of Instr/Cntrl air							
10^{-2} to 10^{-3} / yr	SGTR Stuck open PORV/SV MFLB MSLB inside Loss of 1 SR DC bus				X			
10^{-3} to 10^{-4} / yr	Small LOCA Loss of SW				X			
10^{-4} to 10^{-5} / yr	Medium LOCA Large LOCA (BWR)					X		
<10^{-5} / yr	Large LOCA (PWR) ISLOCA Vessel Rupture						X	

Credit for auto actuation part of the time and operator action to initiate either of two methods of core cooling

Credit for auto actuation part of the time and operator action

Credit for auto actuation part of the time or possibly operator action

1E Non-Software Failure Mechanisms

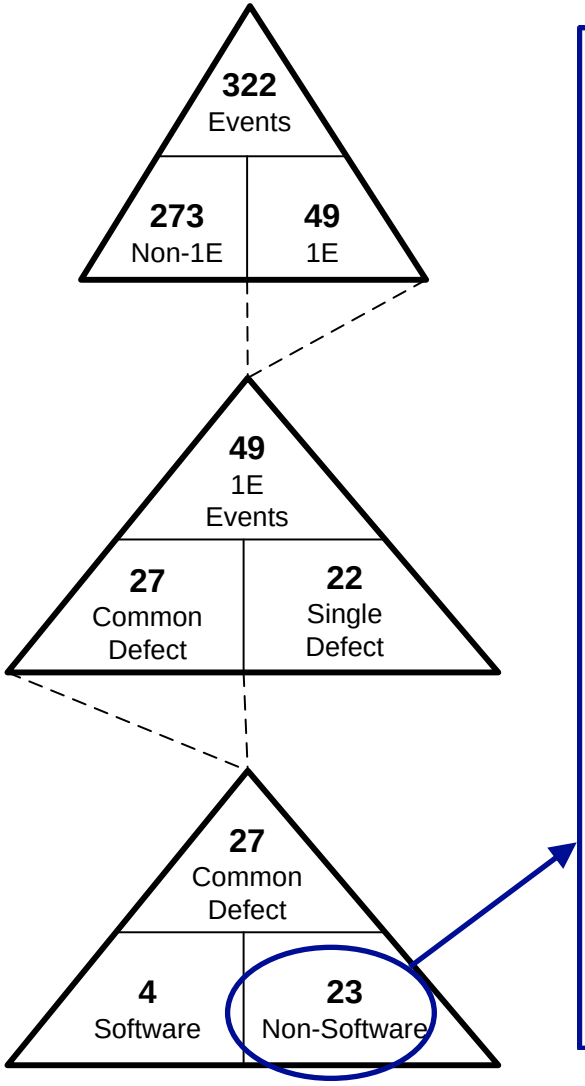

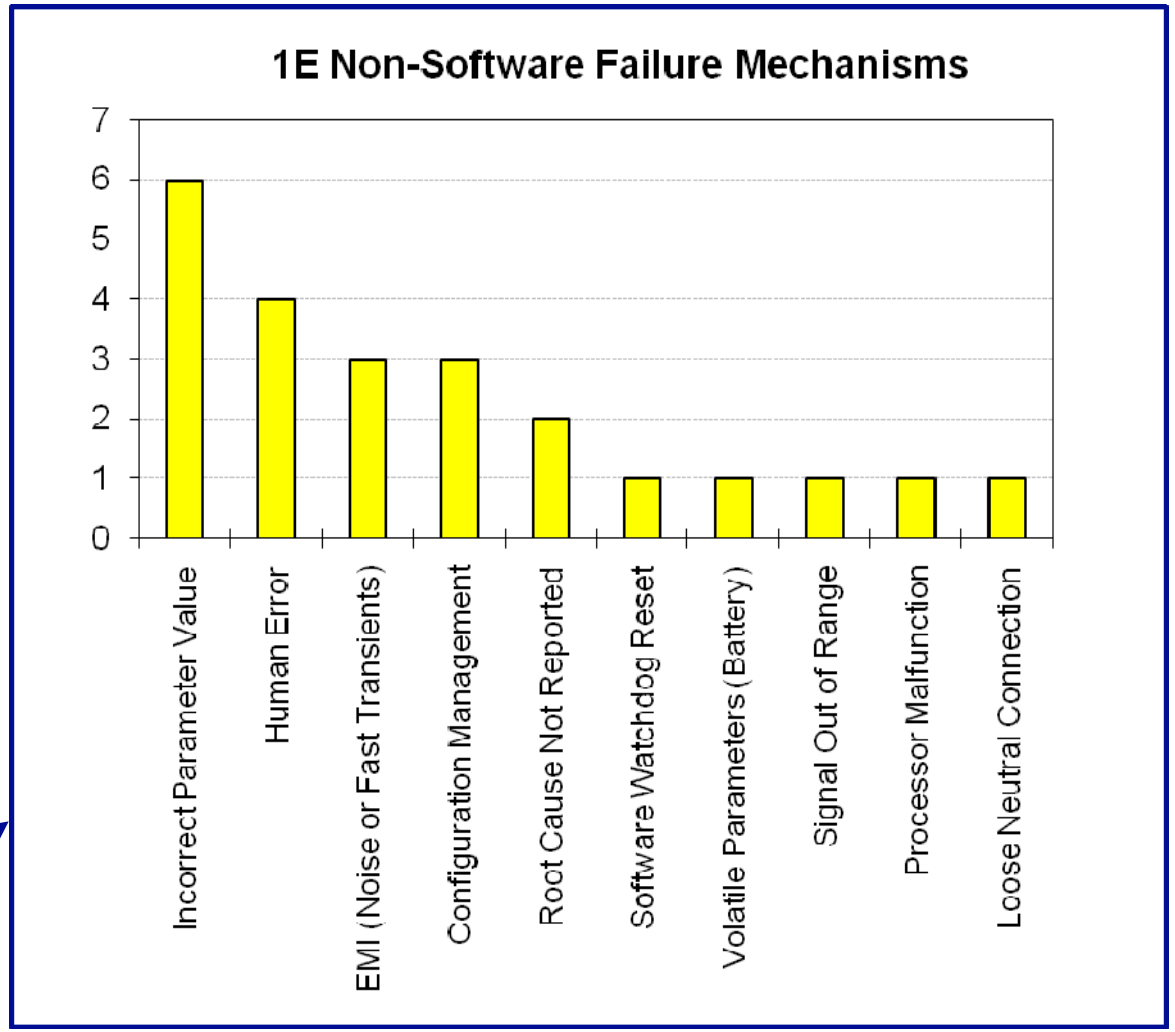

Non-1E Common Defect Events

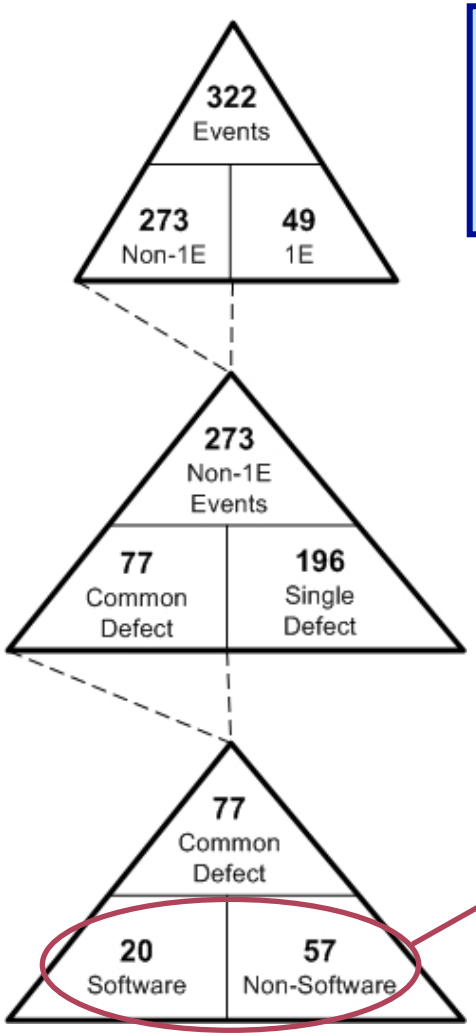

Out of 77 common defect events, 7 resulted in CCFs due to software (9.1%)

		1	1
		2	2
Common Defect?		X	X
Concurrent Triggers?		X	X
Failure?			X
		Potential CCF	Actual CCF
System Level	Software	0	7
	Other	5	26
		5	33

Remaining Common Defects:

- 10 Single Failures
- 12 Spurious Actuations
- 6 Subsystem Level Potential CCFs
- 11 Subsystem Level Actual CCFs

EPRI ELECTRIC POWER RESEARCH INSTITUTE

Non-1E Software Failure Mechanisms

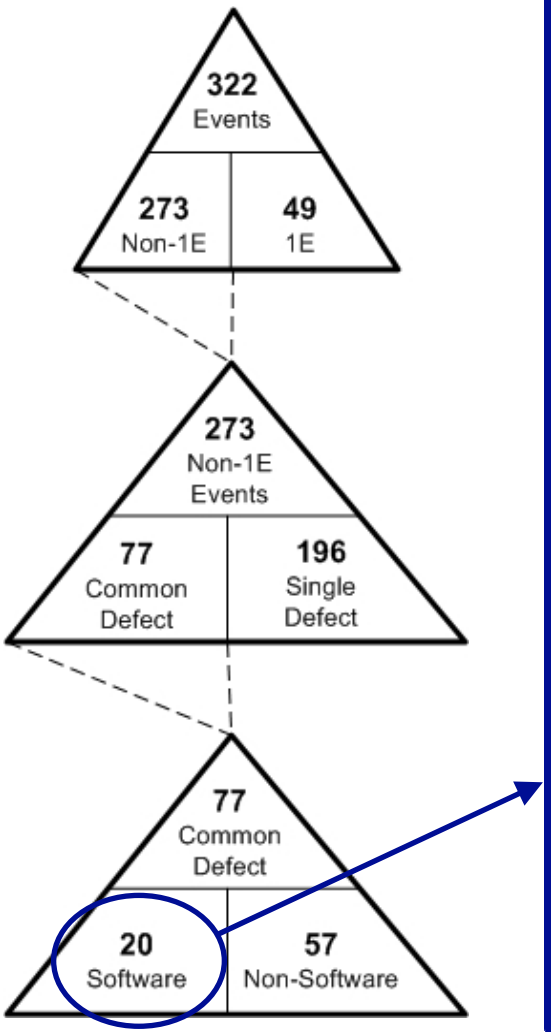

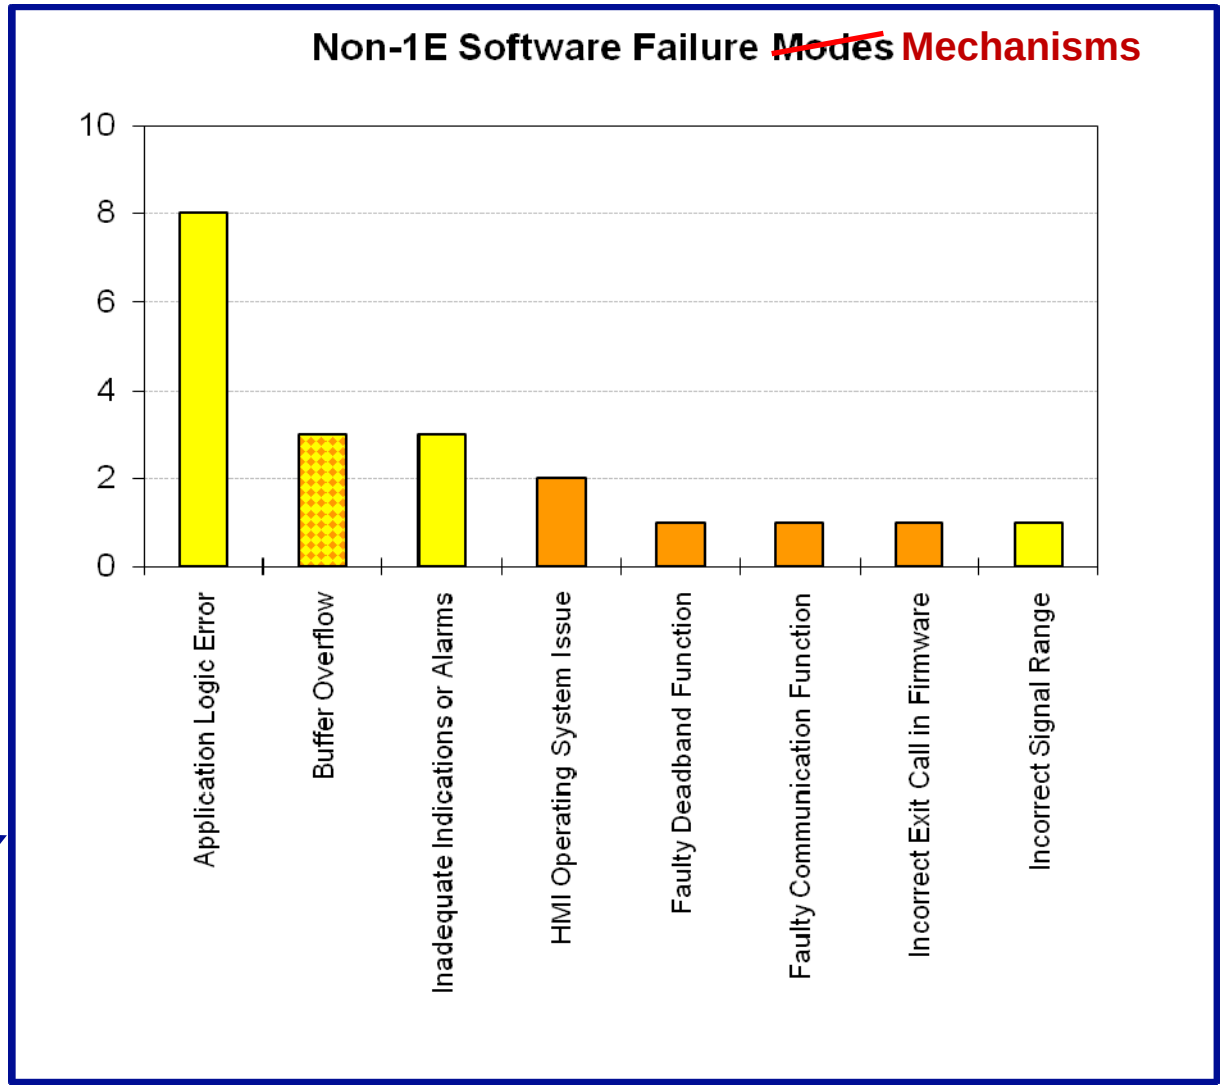

1E vs. Non-1E Vulnerability to CCF

(The System – Not Just The Software)

Attribute	1E Systems	Non-1E Systems
Redundancy	Independent Channels	Master/Slave
Shared Resources	Never	Almost Always
Formal SQA* Methods	Always	Varies (Improving)
Functional Complexity	Low	High
System Interactions	Low	High

Common Defect Events (System Level)	1E Systems	Non-1E Systems
Actual & Potential CCFs:	6 out of 27 (22%)	38 out of 77 (49%)

1E systems are inherently better protected against CCF in the presence of a common defect

*Software Quality Assurance

OE Conclusions

Insights and Inferences

- **Software has been no more problematic than other CCF contributors**
  - Current methods have been effective in keeping software a minor contributor to potential 1E CCFs
- **Difficult to combine 1E and non-1E experience**
- **No events where diverse platforms would have been effective in protecting against CCF**
- **Several events confirmed effectiveness of signal and functional diversity in protecting against CCF**

OE Conclusions, cont'd

Recommendations

- **Capture and promote process and design characteristics that have been effective in protecting against CCF**
- **Encourage additional OE investigations**
  - Other countries and industries (confirm U.S. results)
  - Analyze for
    - Prevalent causes of failures
    - Corrective actions / defensive measures
    - Risk significance

Next Topic:
Digital Failures - Mechanisms, Modes and Effects

- ***"Digital I&C may introduce new failure modes that are not well understood."*** – Letter, Chairman ACRS to Chairman U.S. Nuclear Regulatory Commission, April 29, 2008
- ***"Failure mechanisms produce failure modes which, in turn, have certain effects on system operation."*** - NUREG 0492 (Fault Tree Handbook)
- Discuss:
  - Digital system FMEAs performed today
  - Realistic digital system behaviors
  - Context of nuclear plant safety system
  - Implications for PRA

Digital FMEA Practice

- **Postulate single failures (IEEE 379), follow guidance in IEEE 352**
  - **Tabulate functions, failure mechanisms / modes, channel effects, methods of detection, system effects, remarks & other effects**
- **Deterministic, down to the component level**
  - **Sensors, power supplies, I/O modules, comm. modules, processors, etc.**
  - **Fail high, fail low, fail as-is, loss of comm's, stopped processor, etc.**
- **FMEAs for full 1E upgrades are extensive, 1000+ pages (e.g., Oconee)**
- **Software functions are credited for fault detection and tolerance**

Failure mechanisms and modes are well understood

Digital FMEA Experience

- Helps identify vulnerabilities, protective features
- OE shows mistakes in FMEA can overlook system defects
- Good practices:
  - Use validation tests to confirm expected responses to failure modes/mechanisms, especially methods of detection (e.g., alarms)
  - Use validated FMEAs to help understand & troubleshoot incorrect system behaviors
- FMEAs for full-scale 1E upgrades can be complex and expensive if not managed carefully
  - More efficient treatment may be appropriate
  - Consideration of mechanisms, modes and effects

Failure Mode

Behavior of a system, subsystem or component (viewed from outside) when it fails

- Possible failure modes - determined based on functional requirements, e.g.,
  - For a simple ‘on-off’ protection function:
    - Failure to actuate
    - Late actuation
    - Spurious actuation
  - **Digital system often has same set of possible failure modes as a functionally equivalent analog system**
- Design measures may be used to ensure that particular failure modes are impossible or highly unlikely, e.g.,
  - Cyclic behavior and ‘watchdog’ could rule out late actuation or failure to actuate

Failure Mechanism

An event or chain of events occurring during operation and leading to system or component failure

- Example: A division by zero causes the microprocessor to “crash and freeze”

- Different failure mechanisms could result in the same failure mode
  - Example: A random hardware error or a division by zero could each lead to a spurious actuation
- Design measures can also be used to rule out specific failure mechanisms
  - Example: Absence of divisions (or limiting the denominators) in the executable code ensures that no division by zero will occur

Failure Effect

Impact of a failure mode on the larger component, sub-system or plant system

- Example:
  - Failure mode: One CPC channel does not trip when required
  - Failure effect: No effect at system level - other CPC channels and trip functions scram the reactor
- Ultimately, the failure effects at the safety system and plant levels determine safety

Recall ACRS Letter on Digital Failure Modes *

"Application-independent processor failure modes"
or mechanisms?

- Task Crash
- Task Hang
- Task Late Response
- Task Early Response
- Task Incorrect Response
- Task No Response
- Processor Crash
- Corrupted Input
- Corrupted Output
- Out of Sequence Data

*Letter to Chairman of NRC Commissioners, 4/29/08

Example: Design of a Typical Reactor Protection Function

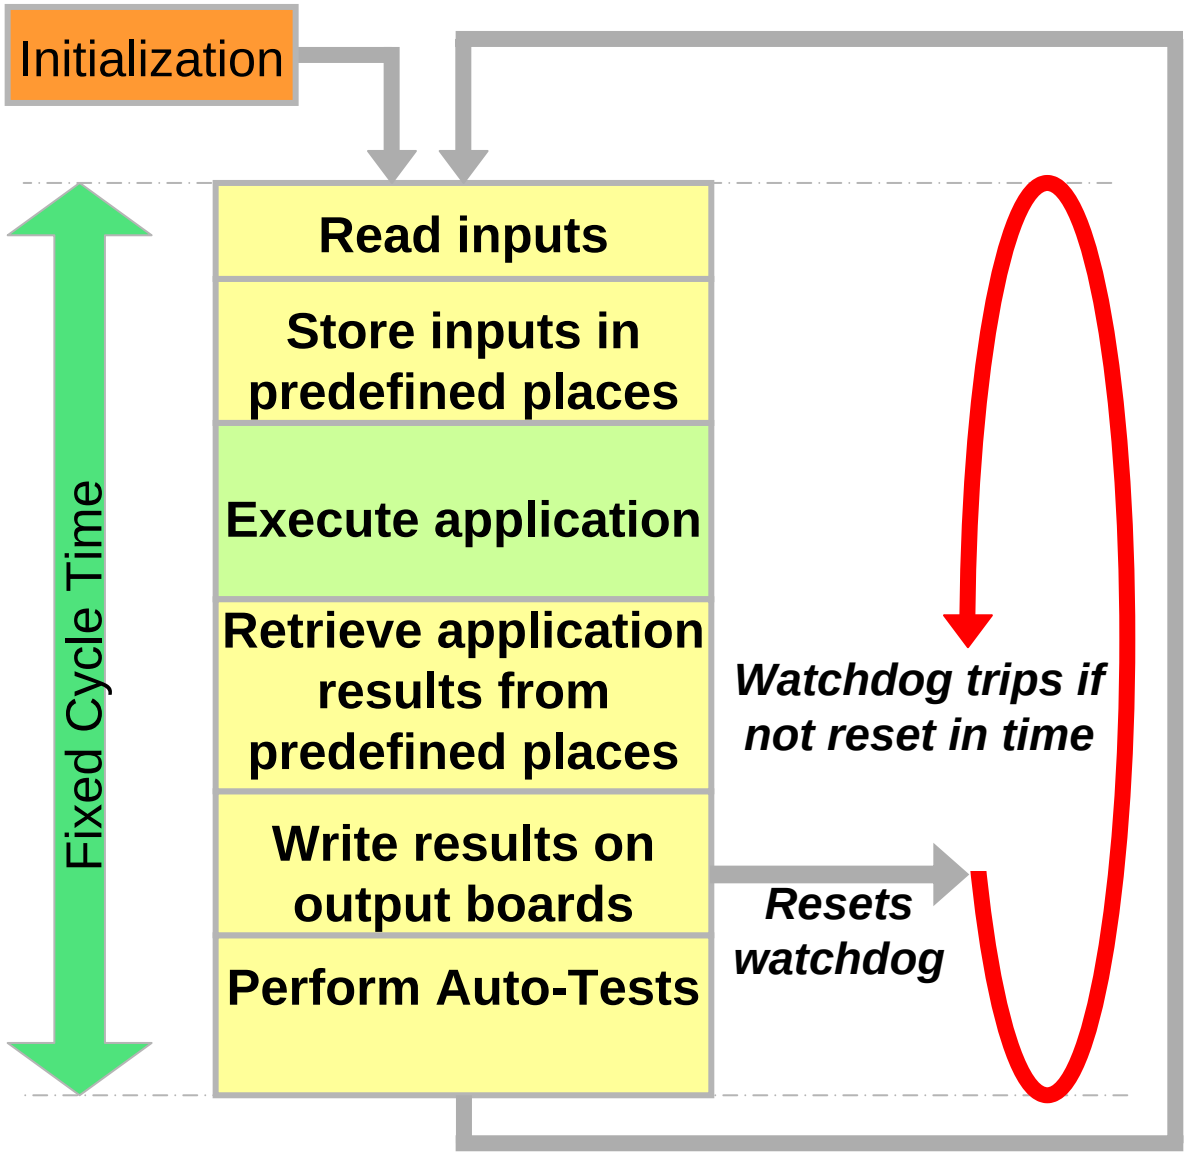

Functions are simple

Outputs are Boolean

1E Systems – Designed for High Reliability

Failure ~~Modes~~/Mechanisms*	Realistic 1E System Behaviors
1. Task Crash	Defensive measure - Any software or processor problem that prevents an output from being issued within a given time frame will cause the hardware watchdog to raise a trip/alarm signal
2. Task Hang	
3. Task Late Response	
4. Task Early Response	
5. Task Incorrect Response	
6. Task No Response	
7. Processor Crash	
8. Corrupted Input	Most digital reactor protection functions use only instantaneous values. Time-dependent functions addressed through programming practices.
9. Corrupted Output	
10. Out-of-Sequence Data	

* From ACRS letter to Chairman of NRC Commissioners, 4/29/08

1E Systems – Designed for High Reliability

Failure Modes/Mechanisms	Realistic 1E System Behaviors
1. Task Crash	
2. Task Hang	**Early responses are not an issue for protection functions. In the worst case, they constitute spurious actuations**
3. Task Late Response	
4. Task Early Response	
5. Task Incorrect Response	**Needs a close look**
6. Task No Response	
7. Processor Crash	**Not specific to digital systems - addressed through redundancy and independence requirements.** **Digital systems usually offer better protection against corrupted inputs than analog systems.**
8. Corrupted Input	
9. Corrupted Output	
10. Out-of-Sequence Data	

* From ACRS letter to Chairman of NRC Commissioners, 4/29/08

Importance of Understanding Digital System Design - Mechanisms for "Task Incorrect Response"

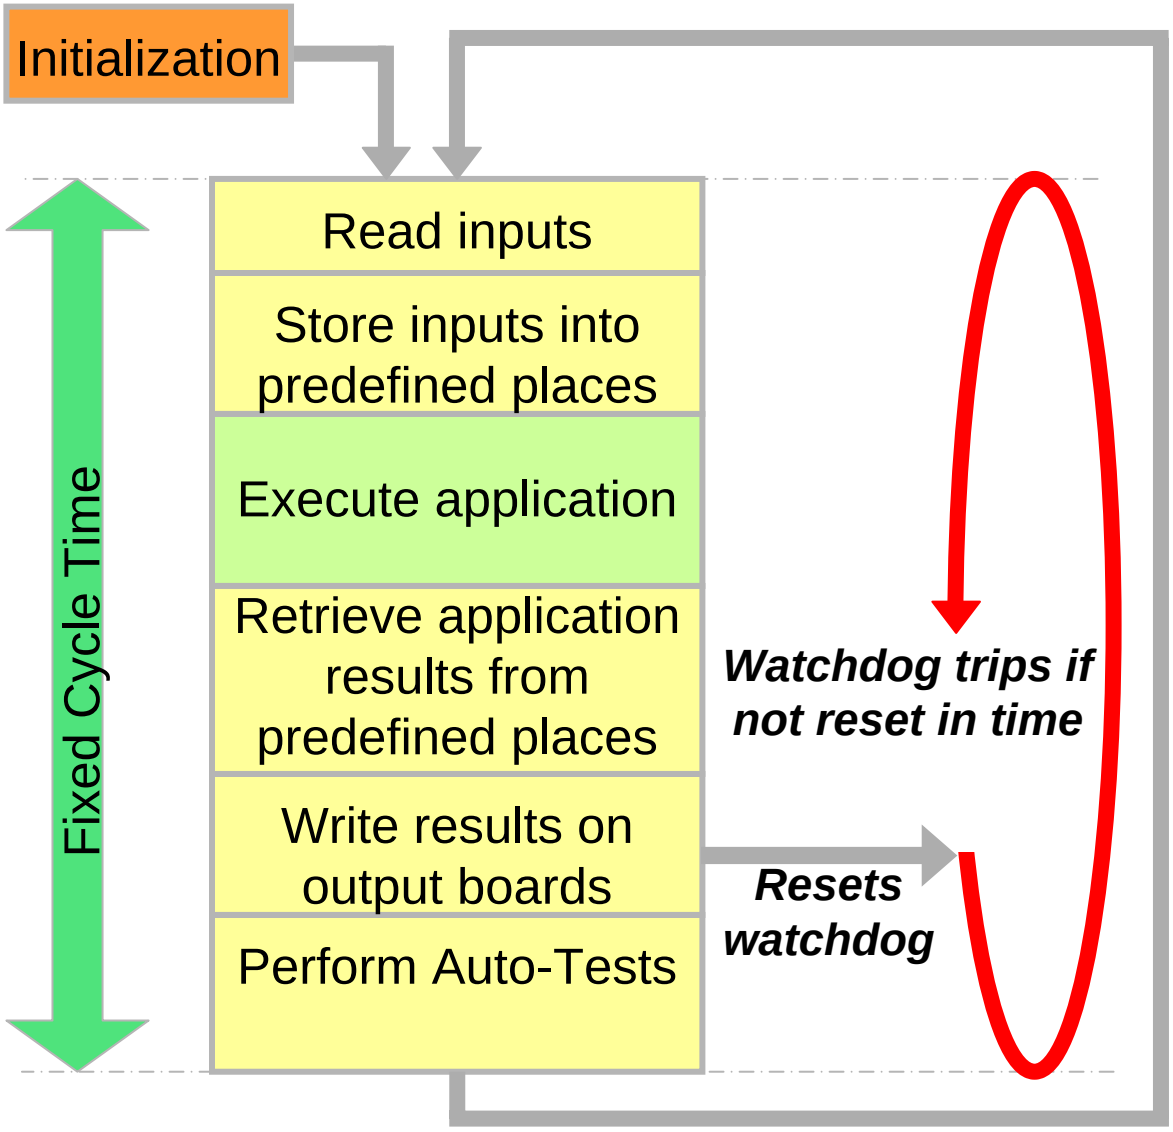

- Operating system fails to read inputs correctly
- Operating system fails to correctly store inputs
- **Application fails to calculate correct results (e.g., OE events 1, 221)**
- Operating system fails to retrieve the correct application results
- Operating system fails to write the retrieved application results correctly
- **Incorrect auto-tests (e.g., OE event #10)**

Digital Failure Mechanisms / Modes / Effects
CCF Implications

- Failure mechanisms may be prevented or mitigated by defensive measures (and/or diversity)
  - Defensive measures can eliminate whole classes of failure mechanisms
  - Not all possible failure mechanisms need to be analyzed - top-down FMEA approach useful
  - Need to identify problematic failure mechanisms
- Diversity may be appropriate, but….
  - Not the only solution, may not be the preferred solution
  - Necessarily adds complexity, but not necessarily safety
  - May be more appropriate between different lines of defense than within a single line of defense
- Match solution to context
  - Integrate diversity, defensive measures and OE insights for CCF protection (prevention as well as mitigation)
  - OE shows importance of defensive measures
  - OE shows benefits of some types of diversity, e.g., functional and signal diversity

Digital Failure Mechanisms / Modes / Effects
PRA Implications

- Only failure modes (not mechanisms) need be represented explicitly in PRA, on an application-specific basis, with:
  - Probabilities of failure modes on demand
  - Frequencies of failure modes in continuous conditions (e.g., spurious actuation and mission time failures, if applicable)
  - Understanding of dominant failure mechanisms may be helpful in estimating failure probabilities and beta factors
- Design measures may prevent or mitigate particular failure mechanisms
  - Good design rules for digital safety systems have been honed and tried and tested over more than three decades
- Some design measures may be effective against a wide range of failure mechanisms
  - Example: Watchdog in earlier slide

Failure Mechanisms, Modes and Effects Conclusions

- Failure **modes** of digital protection systems are well understood
  - System-level behaviors
  - Essentially same as for analog systems
    - Digital system CCF accounted for in D3 coping analysis
    - CCF effects are modeled in PRA for existing plants
  - Extensive FMEAs are being performed by equipment suppliers and licensees
- Failure mechanism evaluation useful to improve the design through incorporation of defenses against problematic failure mechanisms

Failure Mechanisms, Modes and Effects Recommendations

- Future work should continue to develop practical understanding / treatment of digital failure mechanisms, modes and effects, e.g.,
  - Deterministic criteria for defensive measures needed to establish reasonable assurance of adequate protection against classes of failure mechanisms
  - Defensive measure evaluation as basis of reliability estimates for PRA
  - Guidance on options for modeling digital failure mechanisms, modes and effects in PRA
- EPRI and NRC coordinate efforts to develop guidance on protecting against CCF, including complementary use of diversity and defensive measures

Final Topic
Diverse Actuation System (DAS) / Risk Insights

- White paper version transmitted to NRC through NEI in May 2008
  - Risk-informed look at potential benefits and risks associated with automated DAS per ISG-2 (September 2007 version)
- Final EPRI report (1016721) published December 2008
  - Provided to ACRS and NRC January 2009
  - **White paper methodologies and conclusions unchanged**
  - Restructured to improve readability
    - Details moved to appendices, especially sensitivity studies
  - Verbal comments from NRC task working group (TWG) meetings addressed
  - Additional sensitivity study
    - Relative benefits of prevention versus mitigation (suggested by NRC staff)

Diverse Actuation System (DAS)
Discussion Topics

Example of development of risk insights for digital systems using existing PRA methods

- Analysis approach using automated DAS example
  - Deterministic evaluations to identify sequences that might need automated DAS
  - Probabilistic results to assess potential risks/benefits
    - Estimating digital failure probabilities, beta factors
    - **Modeling of failure modes and effects**
- Summarize key insights and conclusions
- Sensitivity studies and effects on conclusions
- Use of risk insights to improve automated DAS design
- Potential impact of revised 30 minute criterion

Application of PRA to Digital I&C Issues
Key Points

- Possible to generate useful risk insights using existing PRA techniques, even without precise knowledge of failure modes and probabilities at the component level
  - **Level of modeling detail commensurate with application**
  - **Results often insensitive to bounding assumptions on failure modes and wide variations in assumed failure probabilities**
- "Conservative" treatment of an individual component or subsystem is not guaranteed to have a conservative impact on the overall system
- Industry developing methods and applying PRA insights to design:
  - Digital systems for new plants
  - Digital upgrades for current plants
- Clarify ACRS statements on application of risk methods to digital

Application of PRA to Digital I&C Issues
Example: Diverse Actuation System for CCF

- Starting point for analysis
  - Analyze each design basis event assuming a coincident software CCF in RPS/ESFAS *
  - Limit credit for operator action as diverse back up to time frames > 30min **
  - Provide an automated diverse actuation system (DAS) for time frames < 30min **

Objective: Demonstrate use of risk methods using automated DAS as an example

* Guidance from Branch Technical Position BTP-19

** Guidance from D3 Task Work Group DI&C-ISG-02. A new revision of the 30 minute criterion has been issued.

PRA Example on Diverse Actuation System – Deterministic Evaluations

- **Purpose: Identify which transient/accident sequences would need an automated DAS per ISG 2**
  - PWRs
    - Westinghouse 2 loop
    - Combustion Engineering
    - Babcock and Wilcox
  - BWR
    - BWR 3
- **Scope of evaluations**
  - Transients
    - Inventory losses at decay heat levels
    - Additional random failures (e.g., stuck open SRV)
  - LOCAs (full spectrum of breaks from small LOCA to double ended guillotine rupture)
  - Steam line breaks (inside outside containment)
  - ATWS

Example Results from Thermal Hydraulic Analyses – Determination of Need for Automated DAS - LOCAs

Case	Purpose	Results	Implications for Automated DAS
PWR LOCA	a. Establish large LOCA range that can be mitigated by low pressure injection b. Determine time to core damage without low pressure injection	Westinghouse ≥ 4” dia CE ≥ 4” dia B&W ≥ 4.5” dia Westinghouse – 2hr CE – 4hr B&W – 45min	To meet the ISG, automated DAS is needed only for low pressure injection
BWR LOCA	a. Establish large LOCA range that can be mitigated by low pressure injection b. Determine time to core damage without low pressure injection	BWR 3 ≥ 4.8” dia < 15 min	To meet the ISG, both high and low pressure injection need automated DAS

Example Results from Thermal Hydraulic Analyses – Determination of Need for Automated DAS – SLB

Case	Purpose	Results	Implications for Automated DAS
PWR steam line breaks	Establish primary coolant system, fuel conditions with: a. ESFAS successful b. No safety injection actuation c. No steam line isolation d. No safety injection actuation or steam line isolation	All fuel temperatures and primary coolant system conditions more benign without ESFAS actuation.	Automated DAS not needed for steam line breaks
BWR steam line break outside containment	Determine time to fuel damage assuming no MSIV closure but with condensate operation	3 hr	Automated DAS not needed for steam line breaks outside containment.

PRA for Digital I&C Example
Deterministic Analyses and Assumptions

- Determine what should be actuated
  - Low pressure injection for PWRs
  - High and low pressure injection for BWRs
- Select conditions that should actuate DAS
  - Multiple diverse indications of relevant accident sequence
    - Low pressurizer pressure **and** high containment pressure in PWRs
    - Low reactor level **and** high drywell pressure in BWRs
  - Require power to actuate (does not actuate on loss of power)
  - Single failure cannot cause spurious actuation

PRA for Digital I&C Example
Probabilistic Analyses

- Evaluate potential benefits of automated DAS
  - Reduction in CDF
  - Reduction in release frequency and offsite consequences
  - Value/impact
- Evaluate potential risks
  - Increase in CDF resulting from inadvertent actuation
  - Compare to benefits
- Perform evaluations for a variety of plant designs
  - 5 PWRs (Westinghouse, CE, B&W)
  - 5 BWRs (BWR 2-6)
- Document risk insights
- Perform sensitivity studies/uncertainty analyses

Benefits/Risks of Automated DAS
BWR Results – Automated DAS for ECCS & MSIV Isolation

Events where DAS is credited

IE	IE Frequency NUREG/CR-1829	Time to 2200°F	P_{OP}	BWR 2	BWR 3	BWR 4	BWR 5	BWR 6
				CDF resulting from digital CCF (P_{ESFAS} ~ 1E-4/dem)				
Large LOCA	2.3E-05yr	<30m		2.3E-09/yr				
Sm/Med LOCA	6.0E-04/yr	>30m	4E-3	2.4E-11/yr				
Med/Large SLB outside cont	*1.0E-04/yr*	*>30m*	*4E-3*	*4E-11/yr*				
Total reduction in CDF due to automated DAS				**2.3E-09/yr**				

Benefits

	Offsite Consequences				
Conditional Containment Failure Probability	0.15	0.02	0.21	0.22	0.01
Person Rem (Large Early Release)	1.5E+06	3.0E+05	6.5E+05	2.5E+06	8.4E+05
Reduction in Dose (person-rem/yr) due to automated DAS	5.63E-04	1.50E-05	3.41E-04	1.38E-03	2.10E-05
Present Value ($2k/person-rem, 3% annual discount rate)	**$17**	**$0.5**	**$10**	**$40**	**$0.6**

		CCDP by plant type				
MSIV Closure		2.6E-06	3.9E-06	6.0E-06	1.4E-06	1.8E-06
General Trans		7.0E-07	1.1E-06	1.6E-06	7.0E-07	7.6E-07

Risks

	IE Frequency NUREG/CR-6928 LERs	Spurious DAS CDF (per year) by plant type				
Spurious MSIV	*0.0024/year*	*6.2E-09*	*9.4E-09*	*1.4E-08*	*3.4E-09*	*4.3E-09*
Spurious Rx Trip	0.0024/year	1.7E-9	2.6E-09	3.8E-09	1.7E-09	1.8E-09
Total increase in CDF due to the automated DAS		***7.9E-9***	***1.2E-08***	***1.8E-08***	***5.0E-09***	***6.1E-09***

EPRI ELECTRIC POWER RESEARCH INSTITUTE

PRA for Digital I&C Example
Level of Detail in Digital I&C Modeling

- Can be useful to model protection system hardware and then assign potential software failure modes to associated hardware.
  - Sensors, communications, voting logic, actuation devices
- I&C also can be modeled at a higher functional level
  - I&C behaviors expressed in terms of behaviors of the components that they control

In our investigation, digital ESFAS was modeled as a super-component

- Similar to the RTS in many current PRAs
- 'The scope, level of detail and technical acceptability of the PRA are to be commensurate with the application…' *RG 1.174*

PRA for Digital I&C Example
Incorporation of Digital System Failure Modes

With I&C design details available

- Develop list of digital components and failure modes from the detailed FMEA
  - *See 'Digital FMEA Practice/Experience' slide*
  - *Software CCF (from D3 evaluation)*

- Without I&C design details
  - Consider failure modes of components controlled by digital system as surrogate failure modes for the I&C system

- **In our investigation, the failure modes of the I&C System were assumed to lead directly to the failure modes of the components they actuated, e.g.,**
  - Pump breakers fail to close
  - MOVs fail to open

Generation of Risk Insights
Estimation of Digital System Failure Probability

- Inputs to failure probability estimate
  - Vendor operating experience
  - Expert opinion based on presence/absence of defensive design measures
  - **International standards, e.g., IEC 60880 (software) and IEC 60987 (hardware)**
    - "For an individual system which incorporates software developed in accordance with the highest quality criteria (IEC 60880 and IEC 60987), a figure of the order of 10^{-4} failure / demand may be an appropriate limit to place on the reliability that may be claimed." Ref IEC 61226
- **In our investigation, initial failure probability assigned assuming high quality design processes - sensitivity studies performed on assumptions for:**
  - **Failure modes**
  - **Failure probabilities**

Generation of Risk Insights
Estimating Spurious Actuation Frequency

- Inputs to spurious actuation frequency estimate:
  - Reviewed 20 years of LERs - general transients, loss of feedwater and loss of main condenser
  - Screened out non-applicable events
  - Of roughly four dozen spurious safety system I&C related events, only 7 were applicable to an automated DAS
    - Not all would be applicable to an automated DAS for the purpose of backing up ESFAS, eliminate those trips resulting from
      - Spurious sensor trips
      - Loss of instrument ac/dc
      - Maintenance/testing errors at power
    - ~0.005/yr spurious actuation frequency

Deterministic Insights from Risk Analysis – Magnitude of Potential Automated DAS Benefits

Benefit relatively small - effective defense-in-depth and diversity provided by existing plant features:

- **Prevention** strategy for LOCA and SLB – provided by reactor coolant pressure boundary:
  - Designed in accordance with piping and pressure vessel codes
  - Periodic inspection per Section XI and pressure vessel codes
  - Monitored during operation (Tech Spec leakage detection activities)
- **Mitigation** of LOCA and SLB – provided by highly reliable ESFAS:
  - Design to consensus standards, redundant, independent trains, etc.
  - Rigorous verification and validation
  - Design features that limit potential for I&C failures and CCF
- **Independence** - Initiating events (LOCA and SLB) and mitigating systems (ESFAS) share no common elements
  - LOCA with loss of ESFAS would require independent failures

Deterministic Insights From the Risk Analysis– Potential Negative DAS Impacts

Potential for spurious operation or manually initiated shutdowns

- DAS intended to mitigate initiating events (large/medium LOCA) that are not expected to occur in any plant over the life of the entire fleet
- It could cause an inadvertent shutdown of a plant somewhere in the fleet once every several years

Note key risk insight for design: DAS should be designed to be robust against spurious operation

Benefits/Risks of Automated DAS
Sensitivity Studies/Uncertainty Analyses

- Numerical issues
  - Set LOCA frequencies to 95% upper bound
  - ESFAS failure probability for automated DAS to be risk-beneficial (~.1)
  - Set human error probability to extremes (BWR only)
  - Parametric uncertainty analysis
- Modeling issues
  - Effects of actuating both high and low pressure systems
  - Effects of actuating the automated DAS on either of two signals
- Completeness issues
  - Failure modes for ECCS
  - Scope of events considered compared to Safety Analysis
  - External events and low power/shutdown operation
  - Include cleanup and lost generation costs in value impact analysis
- Scoping studies
  - Compare BTP-19 and ATWS rule scopes
  - Prevention vs mitigation

Application of PRA to Digital I&C Issues
Risk Insights Applied to Design of Automated DAS

Automated DAS Design Characteristic	Comment
Actuation given multiple plant conditions, all required before actuation	*For example,* *Low pressurizer pressure and high containment pressure (PWR)* *Low reactor level and high drywell pressure (BWR)*
Requires power to actuate	*Not actuate on loss of power (similar to ATWS system)*
No LCOs or allowed outage times in the Technical Specifications	*Availability and reliability requirements determined and performance monitored as a part of Maintenance Rule compliance (similar to ATWS system)*

Application of PRA to Digital I&C Issues
Risk Insights Applied to Design of Automated DAS, cont'd

Automated DAS Design Characteristic	Comment
No automated DAS for steam line breaks downstream of the MSIVs	*Significant time available (BWR)* *Less severe reactor coolant and fuel conditions if automated DAS does not actuate (PWR).*
Timers in series with DAS actuation logic set at latest time to initiate system based on best estimate thermal hydraulic analyses	Similar to BWR ADS timer, allows time for operator intervention to inhibit system actuation in the event of spurious operation
Perform best estimate evaluation of operation of engineered safety features without isolation of plant non-critical systems for a period of 30 minutes	Eliminate need to isolate non-critical cooling systems or shedding of loads needed to support plant operation on actuation of the automated DAS

Application of PRA to Automated DAS Issue
Conclusions

- **Possible to generate risk insights using existing PRA techniques**
- **Automated DAS for events analyzed has little or no benefit**
  - Low frequency events due to prevention measures
  - High quality mitigation systems
  - Independence of initiating events and mitigating systems
  - Spurious transients caused by automated DAS could increase overall risk
  - Conclusions insensitive to digital protection system reliability
- **In general, high frequency events benefit more from augmented defense-in-depth and diversity than rare events**

Application of PRA to Automated DAS Issue Recommendations

- Consider results of this research and encourage Staff and industry use of current PRA methods to address digital I&C issues; e.g.,
  - Where results are insensitive to modeling assumptions
  - Licensing actions, e.g., automated DAS for low frequency events
- Consider revising D3 guidance (BTP-19) to address both event frequency and consequences in assessing adequacy of defense-in-depth
  - Allow a graded approach in which solutions and protective measures are proportional to risk
- Promote methods for addressing of digital system issues that:
  - Credit both prevention and mitigation measures in protecting against failures and CCF
- Clarification of previous ACRS statements on use of PRA methods for digital I&C issues would be helpful

Recap of Key Points

Operating Experience (OE)

- Software no more problematic than other CCF contributors
- Need to capture and promote **process** and **design** characteristics that have been effective in protecting against CCFs

Understanding "Digital" failure modes

- "Failure mechanisms produce failure modes which, in turn, have certain effects on system operation" (i.e., failure modes are understandable)
- PRA models represent failure modes/effects, and do not need exhaustive treatment of low level digital failure mechanisms to generate useful insights
- Failure mechanism prevention and mitigation remain very important in designing robust systems (fault avoidance and fault tolerance)

PRA insights

- Risk insights are possible today using existing techniques
- Need to encourage use of PRA given its capabilities and current state of the art

Recap - Request ACRS Concurrence

Staff and Industry should:

- Continue to gather and apply OE lessons on failure causes, corrective actions and preventive measures – develop common definitions for binning and evaluating events
- Develop methods for crediting defensive measures in protecting against failures and CCF (especially where they are better than diversity), and in assessing digital system reliability
- Use current risk methods to address digital I&C issues for both operating and licensing applications where appropriate, e.g., for low frequency events
- Increase technical exchanges to resolve issues more effectively and efficiently (particularly with RES)

Acronyms

- 1E — Safety system
- BTP — Branch Technical Position
- CCF — Common Cause Failure
- D-3 — Diversity & Defense-in-Depth
- DAS — Diverse Actuation System
- DI&C — Digital Instrumentation and Control
- EPRI — Electric Power Research Institute
- INPO — Institute of Nuclear Power Operations
- ISG — Interim Staff Guidance
- LAR — License Amendment Request
- NEI — Nuclear Energy Institute
- Non-1E — Non-safety system
- OE — Operating Experience
- SQA — Software Quality Assurance
- TWG — Task Working Group

Together…Shaping the Future of Electricity