Apostolakis: On PRA

The chairman of the ACRS's PRA subcommittee offers a clear explanation of the specialized methodology that can help ensure plant safety.

Probabilistic Risk Assessment (PRA) is the systematic process that can be used to examine how nuclear power plant employees and engineered systems work together to ensure plant safety. PRA is quantitative, in that probabilities of events with potential public health consequences are calculated, as are the magnitudes of these potential health consequences. The risk of such events is the product of the event probabilities and their consequences. As practiced in the field of nuclear power, PRA generally focuses on accidents that can severely damage the plant's reactor core and can also challenge the surrounding containment structures, since these pose the greatest potential risk to the public.

PRA integrates into a uniform assessment tool the relevant information about plant design, operational practices, operating history, component reliability, human performance, the physical progression of core-damage accidents, and the potential environmental and health consequences in as realistic a manner as practical.

George Apostolakis is an expert in PRA and a professor of nuclear engineering at the Massachusetts Institute of Technology. He is vice chairman of the Nuclear Regulatory Commission's Advisory Committee on Reactor Safeguards and chairman of its PRA and Human Factors subcommittees. He received the 1999 Tommy Thompson Award from the American Nuclear Society's Nuclear Installations Safety Division.

Apostolakis talked with Rick Michal, *NN* senior associate editor, about how the use of PRA has evolved in the nuclear industry and how it is changing the way the industry is regulated.

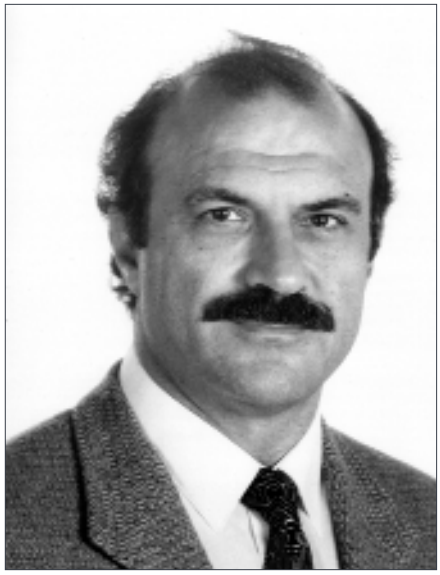

Apostolakis: Use of PRA "raises the safety culture of the plant . . ."

What is PRA?

Probabilistic Risk Assessment, PRA, is an integrated safety analysis methodology that can be summarized by the following four steps:

First, a number of undesirable events must be defined. Some examples of such events are core damage, the release of radioactivity from the containment, and public consequences [fatalities].

Second, for each of these events, the methodology systematically identifies the accident sequences, also called scenarios, that can lead up to it. This is done using logic diagrams, especially for core damage events, that take into account the failure of various emergency safety systems that have been designed for the plant. Phenomenological models are also used as needed, especially for containment phenomena. These scenarios include events such as equipment failures, human errors during tests and maintenance, human recovery actions, loss of coolant accidents [LOCAs], transients, and various external events such as earthquakes and fires. Precisely because these diverse events are included in the accident scenarios, this methodology is called an integrated approach to reactor safety.

Third, the probability of occurrence of each scenario is calculated; thus probability theory is needed. Statistical evidence that is available regarding failures must be taken into account. Also, expert judgments must be used, because many times these events are very rare and have never been seen before.

The fourth and final step, after the probabilities have been identified, is to rank the accident sequences according to their probability of occurrence. This is done because risk must be managed; knowing the major contributors to each undesirable event that was defined in the first step is a major element of

risk management. Also ranked are the SSCs—systems, structures, and components—according to their contribution to the undesirable event.

The ultimate result of the PRA is the probability of each undesirable event and a list of the major contributors to its occurrence.

Is PRA the proper terminology, or should it be PSA, for Probabilistic Safety Assessment?

It depends on the undesirable event. If risk is analyzed—in other words, the undesirable events are latent fatalities or acute fatalities—then I think the proper name is PRA.

On the other hand, if only core damage events or containment failures are analyzed, then PSA is more appropriate.

PRA is primarily used in the United States. In other countries most people use PSA, although now the terms are being used interchangeably.

What is the history of PRA and nuclear power? Have there been successive generations of PRA and has its use in U.S. nuclear power increased over time?

These are interesting questions, because the history of PRA is tied to the history of reactor safety. During the early days of nuclear power development, the 1950s and '60s, people who were designing these facilities and were concerned about licensing them realized that the consequences of a nuclear accident could be catastrophic. Therefore, it was important to keep the probability of these accidents very low. But even though people wanted these probabilities to be very low, they did not have the means for quantifying them. Over the years, a design philosophy evolved that had as cornerstones the concepts of defense in depth and safety margins.

"... [T]he history of PRA is tied to the history of reactor safety."

Defense in depth is usually understood to mean the existence of multiple barriers to prevent the release of radioactivity. Examples of barriers are fuel cladding, primary system pressure boundaries, and the containment structure. Some people expand this definition to mean essentially any measure that enhances our confidence that the probability of accidents is kept low. This would include emergency planning, for example. At the same time, the concept of single failure criterion was developed and a distinction was made between credible and incredible accidents. The focus was on credible accidents. A major focus was protection against large LOCAs. All of this was the deterministic approach to reactor safety.

The first call for a "new approach" to reactor safety was by Reg Farmer, of the U.K. Atomic Energy Authority, in Vienna in 1967. Farmer argued that it was not logical to distinguish between credible and incredible accidents, but that the whole spectrum of accidents should be studied. He used as a measure of risk the release of iodine-131. He proposed to look at sequences of events—the accident scenarios that I mentioned earlier—that lead to release of various amounts of iodine. He also proposed acceptance criteria. Essentially, he formulated the basic idea of PRA.

The first real PRA—the way it is understood now—was published in the United States in 1974 in draft form, and in final form in 1975. It is known as the Reactor Safety Study [WASH-1400], or the Rasmussen report, because Norm Rasmussen, who was a professor at MIT, was the director of that study. The findings of the study created a new thinking about reactor safety. The main concern until that time had been to protect against large LOCAs. The Reactor Safety Study identified as dominant contributors to core damage small LOCAs and transients. The probability of core damage had not been quantified until that time. The Reactor Safety Study came up with numbers for that probability—the best estimate was about five core damage events every 100 000 reactor years—that surprised some people, because they had thought that it was much smaller than this value. The study did an uncertainty analysis and concluded, with very high confidence, that the core damage probability was smaller than three events per 10 000 reactor years, an unexpectedly large number indeed.

"When the Three Mile Island accident occurred in 1979 ... there was renewed interest in PRA methodology."

At the same time, the study showed that the consequences of core damage events were not as significant as previously thought. It also pointed out that operator actions and the support systems, such as the component cooling water system, were very important.

One other important observation is that the Reactor Safety Study identified an important sequence that had been missed until that time. This was the V sequence, which is the failure of two check valves in the PWR emergency core cooling system pressure isolation boundary. This finding demonstrated the value of the integrated approach that I described earlier.

The Reactor Safety Study has an interesting history. It is probably the most reviewed and criticized major study in the history of nuclear power. The focal point of the critics was the executive summary. That summary showed figures that compared the frequencies of certain societal impacts—deaths, for example—from nuclear accidents to those from other man-made phenomena such as airplane crashes. The problem was that the frequency of airplane crashes that killed a certain number of people was known with relatively high precision because it was based on statistics. The frequency of reactor accidents is based primarily on models, judgment, and analysis, and therefore, it is not known as precisely as the other frequency. The critics felt that such comparisons of the frequencies were inappropriate, because the uncertainty about the frequency of nuclear accidents was very large and was not displayed. There was a controversy regarding this point, which unfortunately led the Nuclear Regulatory Commission to decide not to use the Reactor Safety Study in its work. As a result, the study fell victim to the inadequacies of the executive summary.

When the Three Mile Island accident occurred in 1979, people realized that the small LOCA that occurred there was in fact in the Reactor Safety Study. The precise sequence of events that led to the small LOCA was not in the study, of course, but small LOCAs were analyzed in the study. So there was renewed interest in PRA methodology.

The next milestone was the release in 1981 of the PRAs for Zion and Indian Point-2 and -3. The nuclear industry sponsored these PRAs and was aware of the criticisms of the Reactor Safety Study. Extra attention was given to the handling and display of uncertainties in these PRAs. An important result of these studies for Zion and Indian Point was the finding that external events—earthquakes and fires—were significant contributors to risk for these facilities. These events had been dismissed by the Reactor Safety Study. Another major result revealed in these studies was that the containment did not always fail following a severe core damage event.

The next major milestone is the NUREG-1150 study by the NRC that was issued in 1989. This one looked at five plants, and the focus was on severe accidents and containment performance. A general finding of the study was that risks were lower than calculated in the Reactor Safety Study. This was attributed to a better understanding of accidents and better models, because this study was released about 15 years after the Reactor Safety Study.

The NUREG-1150 study used expert judgment to estimate various parameters of phenomena that are expected to occur in the containment after a core damage event. This turned out to be controversial, because objections were raised over the fact that experiments and data were being replaced by expert judgments. But this was not true. The whole idea was to develop a snapshot in time of the risks. There simply was not the time nor the re-

sources to run all the experiments that would be needed to get this harder information.

The NRC issued a generic letter [GL 88-20] in 1988 requesting that each licensee in the United States use PRA-like methodologies to perform a plant-specific search for vulnerabilities that might lead to severe accidents. These studies are known as the Individual Plant Examinations, IPEs. The NRC completed its review of the program a few years ago. This program was successful in the sense that both the NRC staff and nuclear power plant personnel familiarized themselves with the methodology of PRA.

Then a major milestone occurred in 1995, when the NRC reversed itself and issued the PRA policy statement that directed the NRC staff to use PRA in all regulatory matters to the extent supported by the state of the art. However, the Commissioners included an important statement, which was that PRA's use should be in a manner that complemented the defense-in-depth philosophy. This is important because it shows how cautious the NRC was regarding the use of PRA. Defense in depth, the traditional cornerstone of reactor safety, was placed at a higher level than PRA. The Commissioners also encouraged the NRC staff to use PRA to reduce unnecessary conservatism associated with current regulatory requirements.

Following this, a couple of years ago, the NRC issued a number of risk-informed regulatory guides. Regulatory Guide 1.174, issued in July 1998, is a major milestone because it states how PRA can be used formally when a licensee requests a change in the licensing basis. It lists a number of principles and expectations, and goes into detail as to how to do that. In my view, it is one of the major milestones.

Has PRA's specialized nature retarded its acceptance in additional applications in the nuclear industry?

Yes, I think that is true. It goes back to the strength of the PRA. I said earlier that PRA is an integrated approach to reactor safety and I mentioned that in the accident sequences such things as LOCAs and human errors have to be included. Each one of these requires special expertise. The analysis team must include experts in human performance assessment, experts in how equipment works, and so on. This appears overwhelming to some people, and I think it has contributed to the slow spread of the methodology.

I think it is unrealistic to expect the user of PRA—or even the PRA specialist—to be an expert in all of these diverse disciplines. The users—and I think we're going to have many more users than PRA specialists—have to be aware of the major assumptions behind the PRA models, so that potentially erroneous conclusions will be avoided. That is easier said than done. How it is to be accomplished is something that the industry is learning now how to do.

There is another reason why I think PRA has not been as popular as it should be, and that is that over the years, since the days of the Reactor Safety Study, the NRC has been eager to take negative results of PRA—such as accident sequences that had not been identified before—and issue new relevant regulations.

On the other hand, PRA results were never used to relax some of the NRC's existing regulations. So the industry saw PRA as perhaps only an excuse for imposing new regulations. This of course dampened the enthusiasm the industry had for it. But with the new risk-informed regulatory guides that were issued by the NRC a couple of years ago, I am confident that the use of PRA will be much wider in the near future.

How does the use of PRA in U.S. nuclear power compare with similar use in Europe and Asia?

It varies a lot from country to country, especially in its scope. In some countries, there are complete full-scope PRAs that calculate public risk. In others, there are more limited PRAs that stop at the release of radioactivity or core damage. Where the United States really differs is in spending a lot of effort now to risk-inform the regulations. This is something that is being pioneered here. I think the rest of the world is looking to the NRC and the American industry to see how this revolution will materialize.

What are the strengths of PRA in the United States?

I believe the most important strength is the ranking of accident sequences and the ranking of the SSCs, because these rankings are essential to rational risk management and the wise allocation of resources. These results come from analyses that include everything that could be thought of that can go wrong at the facility.

This integrated approach is the new thing that PRA is introducing. Defense in depth was applied earlier to individual issues, individual systems of the plant, as well as the whole plant, without the benefit of the integrated approach of PRA and its ranking of accident sequences and SSCs. The result was that unnecessary regulatory burden was created in some instances, such as in quality assurance requirements, and at the same time some important accident sequences were overlooked, such as the V sequence that I mentioned earlier.

Another very important strength of PRA is its value as a communication tool. The analysts and users can use PRA diagrams that depict the accident sequences to any desired level of detail to communicate to others their work. I have found this to be very useful. A reviewer can now express his or her disagreement in specific technical terms and the ensuing debate is a very healthy step toward consensus.

What about weaknesses that exist in PRA, and how could they be eliminated?

I think that it's clear by now that PRA is very ambitious. It looks at the plant as an integrated system, and because of this ambitious approach there are several areas where improvements are still needed. There are issues of scope; for example, low-power and shutdown modes of operation must be understood better than they are today. There also are modeling needs, such as human reliability assessment. Perhaps there needs to be a better job done in assessing the risk from fires and so on.

Could you give examples of how PRA has made contributions in such areas as on-line maintenance and improving NRC regulations?

I think the main use for PRA has been in evaluating very quickly the core damage frequency for different configurations. For plants that do on-line maintenance, what needs to be understood is the impact on safety of the maintenance tasks that are being performed on line. So PRA provides the tool to determine quickly what could be taken out of service and for how long. It essentially helps control the plant configurations, typically by looking at the changes in the core damage frequency that result from taking certain equipment out of service.

The same thing applies to the NRC Maintenance Rule that was implemented in 1996. The Maintenance Rule requires that the SSCs important to safety be maintained so that they will perform their safety functions when required. PRA, using the ranking methods that I mentioned earlier, reveals which SSCs are risk-significant. It helps to set the performance criteria for these SSCs, such as maximum unavailabilities. It helps to assess the impact of the removal of SSCs from service.

PRA is also used to manage outages. This is an interesting area, because there is still work to be done for low-power and shutdown modes. But there is also work that has already been done. What the utilities are using now to manage outages is a combination of defense-in-depth measures and risk insights from the PRAs that have been done for shutdown modes.

PRA has also been used in training. The idea here is to improve human performance. The operators, by studying the PRA, learn what the dominant accident sequences are, which accident sequences involve critical operator actions, and why. These can then be discussed in groups. So the level of understanding increases. Also, PRA can help select important scenarios to run on simulators.

Regarding risk-informed regulatory applications, the major regulatory guides that were issued a couple of years ago deal with such areas as risk-informed in-service inspection, risk-informed graded quality assurance, risk-informed technical specifications, and risk-informed in-service testing. Some utilities are using these regulatory guides already. One utility has indicated that if it implemented only the graded quality assurance guidance, its savings would be up to $2 million a year, because it now spends a lot of money on quality assurance requirements for SSCs that are insignificant from the risk point of view.

Another major current activity is the NRC's revision of the reactor oversight process, which is becoming risk informed in a number of ways. This is still in the pilot stage, but it represents a major change in the NRC inspection and enforcement process.

Continued

Could you discuss PRA standards?

One of the conclusions of the review of the Individual Plant Examinations was that there was large variability in the results from unit to unit. Part of this can be explained as being due to design differences, but another part is due to different models and methods that the individual utilities used. Given this state of affairs, there is a real issue here. For example, if a utility comes before the NRC using the risk-informed approach to request something, how does the NRC decide that this is a good analysis? So, standards are very desirable. The problem with standards is that in some areas where more development is needed, perhaps progress will be inhibited, because once something is published, people think that it is good enough.

The American Society of Mechanical Engineers is working on a PRA standard for internal events, and the American Nuclear Society is developing another standard for external events—earthquakes, fires, tornadoes—and also for low-power and shutdown modes.

I think it's a good idea to have a standard, especially if the standard specifies minimum requirements for a good PRA. For example, there can be a list of various kinds of LOCAs and transients that must be included in the PRA. Common cause failures must be analyzed. But I'm not sure that a standard should actually specify what is a good PRA, because what is good or adequate depends on the application.

How has PRA been used in other industries?

There are several applications of PRA in other industries. The chemical industry is calling it QRA, Quantitative Risk Assessment, rather than PRA. The Center for Chemical Process Safety has been issuing a series of books that provide guidance on how to do it, and an experienced nuclear PRA practitioner can open these books and recognize that the methods there are very similar to the ones that the nuclear industry is using.

In the aerospace industry, PRA has been applied to the space shuttle and to the Cassini mission RTGs [radioisotope thermoelectric generators].

An interesting application of PRA was sponsored by the U.S. Army. The Army, as part of its chemical stockpile disposal program, has built a facility in Utah, which is called the Tooele chemical agent disposal facility. It is an incinerator that disposes of extremely hazardous chemical agents and munitions. The Army commissioned a complete PRA that estimates public and worker health consequences from the incineration process. They also estimated the risk from continuing storage of the chemical munitions. Thus, the Army had all the information that was needed to make a decision as to whether to go ahead and incinerate these chemical munitions or just keep them in storage. This was really an impressive piece of work, and it was requested by the National Academy of Sciences, which oversees the activities of the Army in this respect.

"One utility has indicated that if it implemented only the graded quality assurance guidance, its savings would be up to $2 million a year. . . ."

Also, PRAs have been used to analyze the hazards from offshore structures such as oil rigs. They are not exactly like the complete methodology that we call PRA for reactors, but the ideas are there and a lot of the tools like fault trees and event trees are used.

How could the use of information technology improve the use of PRA?

I think it's already improving it and has had a strong impact. The main impact is that there is now the ability to calculate very quickly the core damage frequency or the frequency of large releases for different configurations and under various assumptions. Sensitivity studies can be done very quickly. In applications such as on-line maintenance or the Maintenance Rule, that capability is needed to produce results for different configurations very quickly, as I said before. This has been a major change. I remember 15 to 20 years ago when a change in something in the PRA meant perhaps days of calculations to get the final result. Now it can be done in a minute or two.

Utilities are also installing safety monitors. This means that television monitors are installed in various rooms of the facility that show the current core damage frequency given the reactor's configuration. Southern California Edison's safety monitor, for example, solves a complete PRA in less than one minute at its San Onofre Nuclear Generating Station. This is remarkable. An interesting observation here, something that was unexpected: By placing all these television monitors in all these rooms, now the people who work at the plant—not just the operators and the PRA analysts but the general personnel there—can look up and see what the current value of the core damage frequency is. Sometimes they know that if they complete the task they are performing at that time, the core damage frequency will decrease. So it raises the safety culture of the plant, which is an important if unexpected result. NN