D930318

The Honorable Ivan Selin
Chairman
U.S. Nuclear Regulatory Commission
Washington, D.C. 20555

Dear Chairman Selin:

SUBJECT: COMPUTERS IN NUCLEAR POWER PLANT OPERATIONS

During the 395th meeting of the Advisory Committee on Reactor Safeguards, March 11-12, 1993, we discussed the staff's progress in defining the regulatory requirements for digital instrumentation and control systems. During this meeting, we had the benefit of discussions with members of the NRC staff.

We have now had a long series of meetings, and have heard from many relevant people, but by no means all. To some extent our input has been biased in the direction of people, groups, and organizations who have experienced problems, and we have not heard from the legions of organizations who have successfully made the move into the computer world. It is important not to develop a tabloid mentality about new technology, i.e., aberrations from the norm treated as if they were the norm.

A first observation is that many of the anecdotes about catastrophic failures of major computer systems refer to systems far larger than those of interest here. Even the software systems on the C-17 aircraft, written in nearly a dozen languages for nearly a dozen machines, are far larger than any of relevance to the nuclear business. The Strategic Defense Initiative dispute is even less relevant. So we have to maintain perspective about scale.

A second observation is that computerization provides an opportunity, not a threat. The extraordinary reliability of electronic systems (unless abused), their potential for continuous and extensive self-testing in real time, their potential for relatively painless upgrades as experience accumulates, their ability to cover an enormous function space and to accommodate unseemly amounts of input data, their remarkable immunity to wear (few, if any, moving parts)ýall these provide the potential for safety enhancement. Much of our input from the staff has been devoted to the negative aspects of computerization, as if it were a disease to be kept in check.

A related observation is that the transition to computerized operation, control, instrumentation, support, recordkeeping, and maintenance procedures and records, is inevitable. The job of the NRC is not to manage or resist the transition, but to maintain a reasonable level of assurance that it is accomplished with proper accounting for the impact on safety. With any reasonable use of the technology the impact is expected to be large and positive.

The regulatory issues we have isolated in our series of subcommittee meetings fall broadly into two categories. One is a consequence of lack of nuclear regulatory experience with modern electronics, especially computers, leading to both extraordinary conservatism relative to unfamiliar accident sequences, and the application to a new technology of review methods and nomenclature derived from old habit and experience. The second is a collection of genuinely new problems associated both with the complexity of the new technology and with the consequent difficulty of assessing (as distinguished from assuring) its level of safety. We deal with these in order.

Failures of computerized systems (excluding fans, hard disks, and other mechanical components) do not follow the traditional bathtub curve of infant mortality, stable performance, and then wearout. Electrons don't wear out. Both in electronic hardware and software there tends to be a period of infant and young adult mortality (to which we will return), with performance and reliability gradually improving with time simply through natural selectionýbugs are ironed out through experience and through extensive testing. There is no later period of wear, so there is no place for the regulatory and maintenance procedures associated with that part of the reliability pattern. Further, self-testing can provide constant assurance of full functionality of the electronics.

As a consequence, however, there has been little progress in applying the methods of probabilistic risk analysis, on which we have become so heavily dependent for mechanical, hydraulic, and electromechanical systems, to computer systems. Indeed the semi-conductor components of the computerized systems are inherently so reliable that high-temperature life-testing is the only means available, in most cases, for generating any failures at all. Whereas one can generate probabilities for the existence of perinatal defects, there is no such thing as a probability per unit time for the development of disease. Nor does in-service inspection play the same role.

These are important points, because the concepts of reliability and reproducibility differ, and the testing and verification procedures used depend on which is to be assured. A mechanical component with a presumed reliability of 10-3 failures per demand can be tested a few thousand times to assure that level of reliability, but a software-based system with a hidden bug that will be revealed in the event of an unlikely input configuration can be tested without failure until the cows come home, but will still always fail with that particular input. Interest has therefore to be directed at the probability that there is such a hidden bug, and the probability that some other circumstance may generate the unfortunate input. Neither of these probabilities will be discovered by repetitive testing under normal conditions. Randomized input testing can tell one something about the former probability, but not the latter. It is therefore misleading to bandy failure probabilities around, as if they had the same meaning as they do for familiar mechanical and electrical components. It also makes the direct comparison of computerized system reliability with the reliability of older technology more difficult.

These and other considerations mandate a format adjustment for the regulatory system, and such changes tend to be painful. What we have seen here is an unfortunate effort to cling to the old ways, to the point of asking that all digital systems have analog backupsýnot because the latter are better or more reliable, but because they are more familiar to the regulator and therefore easier to regulate. That alone could place an unwarranted burden on those seeking to improve safety by updating technology.

The second category of issues follows from the undoubted fact that computerized systems do indeed introduce unfamiliar failure modes, which require both recognition and palliative measures. Too much attention appears to have been concentrated on a microcosm of the more recognizable of these matters, specifically vulnerability of digital systems to electromagnetic interference (a subject on which there is enormous military expertise, largely untapped by the NRC staff), and the fact that replicated defective software (like replicated defective hardware) can be the source of common-mode failures. Both of these are real issues, but, in our judgment, not the central ones.

Let us first consider software issues. The literature is full of examples of cases in which carefully written and tested software still contains errors. Indeed it is doubtless true, though in principle unprovable, that any large program that has not undergone a formal verification and validation (V&V) contains yet undiscovered errors. Lest there be confusion, it is well to be quantitative about the problem of implementing a function in software.

The simplest of all digital programs might generate a logic function, a mapping that accepts a number of binary inputs (say n) and generates a single binary outputýa signal that might, in turn, activate a pump or a valve or some other sequence of events. Such a logic function has 2n possible input states, over a thousand for n=10 and over a million for n=20. These are not unreasonable numbers of input states, because the input of a single number to one percent accuracy requires seven (usually more) binary inputs. Since each input state can have either output state (on/off), that means that even a modest eight-input binary converter of this sort can represent 2256 or 1077 different logic functions. A defect (either hardware or software) can change the desired function into any of the others. It is therefore reasonable to expect to test the system to make sure that it performs as designed, but not reasonable to expect to explore, by brute force, all consequences of all possible defects. The point is only strengthened if one has more complex outputs than just a single bit.

If, therefore, the requirements specified for the system describe the full mapping of the input space to the output space, special methods will be required to verify that this has been accomplished correctly. Such methods exist, and are applicable to relatively simple software packages. When formal V&V is possible, it provides assurance that the code, as written, correctly implements the formal specifications laid upon the design. When it is not possible (because the code is too long or too complex), there are many alternatives, but none of them provides the kind of assurance

of code fidelity that is provided by formal V&V.

There appears to be a consensus among the experts we have consulted that the safety-related software in nuclear power plants is within reach of formal V&V methods, and that the potential for serious error lies more in incorrect expression of the specifications than in incorrect programming. Formal V&V can assure that the code correctly expresses the specifications, but not that the specifications are correct. In either case, it would appear that the staff emphasis on the possibility of common-mode errors in code segments used in different parts of the instrumentation and control system is misdirected. We continue to see an urgent need for staff augmentation with people experienced in thinking in the terms outlined above.

We believe that the experience of other industries that have accepted the progress has been characterized, almost without exception, by increases in efficiency and reliability, and by concomitant decreases in cost. (While the latter is not the NRC's business, it remains true that resources and attention released from unproductive safety concerns may, at least in part, find their way to better use.) There are genuine safety issues in this transition, of which one unfamiliar one is surely the requirement, in order to generate verifiable software, for precise no-nonsense attention to the specification of the functions to be implemented by the software.

The gist of our concerns is that the regulatory procedures developed during the decades preceding the full flowering of the electronic revolution (which may not yet have occurred) are inappropriate to the regulation of computerized functions in nuclear power plants. (This is true for both hardware and softwareýtoo much emphasis on the distinction is not helpful.) As a consequence, the staff has been dealing with the problems that have shown up so far on an ad hoc basis, applying methods created for each problem, with little underlying methodology. That has resulted in such distractions as the analog-to-digital conversion problem, the overemphasis on electromagnetic interference problems, the singling out of software common-mode failure as a central issue, etc., all without a framework into which the broad issues of regulatory emphasis and consistency can be fitted. We can cavil about the specific staff approaches to each of these, but the central issue is that neither the staff nor the Commission has established what could be described as a standard review plan or even a regulatory guide that could help both the staff and the industry know what is expected of them. A statement of the applicable standards ought to precede, not follow, their application. Without such a definition of objectives, coherence is an inevitable victim.

What, then, do we recommend? We frankly doubt that a coherent and effective review plan for computerized applications in nuclear power plants will be produced by the staff, the Commission (whose job is at a higher policy level), or the Committee (which is limited in both resources and expertise). Still, if one believes (as we do) that it needs to be done, it will be necessary to bring in outside help. It was in that context that we initiated our long

series of subcommittee meetings on the subject. Our recommendation is that a workshop and study (with a charter to produce such a plan) be commissioned to be done by the National Academies of Sciences and Engineering. To derive maximum benefit from such a study, there should be appropriate participation by key senior members of the staff.

Additional comments by ACRS Members James C. Carroll and Carlyle Michelson are presented below.

Sincerely,

Paul Shewmon
Chairman

Additional Comments by ACRS Members James C. Carroll and Carlyle Michelson

We agree with most of the technical observations made in this report. However, we disagree with the report's recommendation that a workshop and study be undertaken by the National Academies of Sciences and Engineering for the purpose of developing a review plan for computerized applications in nuclear power plants. Contrary to the view of our colleagues, we believe that the staff and its consultants are making satisfactory progress toward developing a "coherent and effective" review plan. Ideally, such a plan should have been developed in advance of the receipt of applications for the use of this rapidly changing technology. As a practical matter, it has been necessary for the staff to interact with the first group of applicants proposing computerized systems in order to gain an understanding of these systems. This has been a necessary first step before a generic review plan can be developed. Our view is that the proposed National Academies of Sciences and Engineering workshop and study would add little to the process of developing a staff review plan at this point in time.

We note that the staff has attended the series of ACRS subcommittee meetings on computerized applications in nuclear power plants that form the basis for this Committee report. In addition, the staff is planning to sponsor a workshop this fall and plans to obtain ACRS feedback on speakers and topics to be covered.