AIM

Association for Information and Image Management

1100 Wayne Avenue, Suite 1100
Silver Spring, Maryland 20910
301/587-8202

Centimeter

Inches

MANUFACTURED TO AIM STANDARDS
BY APPLIED IMAGE, INC.

O

UCRL-JC-116038
PREPRINT

EVALUATING SOFTWARE FOR SAFETY SYSTEMS IN NUCLEAR POWER PLANTS

J. Dennis Lawrence
Warren L. Persons
G. Gary Preckshot

This paper was prepared for submittal to the
9th Annual Conference on Computer Assurance
Gaithersburg, MD
June 27-30, 1994

January 11, 1994

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This work was supported by the United States Nuclear Regulatory Commission under a Memorandum of Understanding with the United States Department of Energy, and performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

Evaluating Software for Safety Systems in Nuclear Power Plants

**J. Dennis Lawrence
Warren L. Persons
G. Gary Preckshot**

January 1994

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Evaluating Software for Safety Systems in Nuclear Power Plants

**J. Dennis Lawrence, Warren L. Persons and G. Gary Preckshot
Computer Safety & Reliability Group
Lawrence Livermore National Laboratory**

**John Gallagher
Human Instrumentation and Controls Branch
Office of Nuclear Reactor Regulation
U. S. Nuclear Regulatory Commission**

1. Background

Commercial nuclear reactor technology has been in existence for more than forty years. At the end of 1992, 419 commercial reactors were in operation with an additional 59 under construction [1]. (The U.S. numbers are 112 and 6, respectively.) In comparison with other technologies, reactor technology has had an excellent safety record. There is considerable interest in continuing and improving that record.

Existing reactors generally use separate instrumentation and control (I&C) systems for control and for protection — indeed, this is mandated by law in the United States. Both types of I&C systems have historically used analog systems coupled with either electro-mechanical relay logic or hardwired solid state logic as the basic technology.

In the past in the United States, most reactors have been designed individually, built by a company specializing in reactor technology, owned and operated by an electric power utility, and regulated by the U.S. Nuclear Regulatory Commission (NRC). Regulation is limited to assuring public safety to the maximum extent possible. Regulation has included the granting of a construction license before construction of a reactor begins, granting of an operating license after construction but before operation, and inspections during construction and operation.

A number of changes are taking place both in I&C technology and in regulation. New reactor designs will permit reuse of a design at multiple sites. At the same time, I&C technology is moving toward the use of digital computers. Moreover, U.S. law has expanded the number of methods of licensing. In the future, the NRC may issue a design certification for a particular reactor design. When a utility wishes to build a reactor, it will obtain a combined construction and operating license, and may use any certified design. Once built, the plant can begin operation without any additional public hearings, but not until certification inspections, tests and analyses meet certified acceptance criteria.

As part of the design certification, the applicant submits a set of inspections, tests, analyses and acceptance criteria (ITAAC) such that if the inspections, tests and analyses are carried out, and they meet the acceptance criteria, the plant as constructed will be safe. The applicant carries out the ITAAC, and the NRC assures that it is carried out using an audit process. A design certification will be in force for fifteen years. A plant constructed using such a design might well operate for forty years, with the possibility of license renewal for an additional thirty to forty years. The result is that a plant built to a design certified today may still be operating nearly a century from now. Thus, it is imperative that the design be safe!

The safety of the design is not expected to be a problem for many aspects of the nuclear power plant: concrete and steel technology is reasonably well understood, fairly stable, and not likely to change over the next fifteen years. A potential area of concern is the computer-based I&C systems that control the plant and the plant protection systems. The latter is of particular concern to the NRC.

The Lawrence Livermore National Laboratory (LLNL) has been working on aspects of nuclear power technology for nearly half a century, both for the Department of Energy and its predecessors and for the NRC. This includes work on seismic analysis, transportation of nuclear materials, nuclear physics, fusion power, long term waste storage, isotope separation, radiation health issues, digital I&C issues, and the like.

Much of the experimental work in physics conducted at LLNL requires the use of automated control systems, so the Laboratory has been building some of the world's most advanced digital control systems for decades. Examples include the Magnetic Fusion Test Facility, the Nova facility and the Atomic Vapor Laser Isotope Separation facility.

In 1991, LLNL was asked by the NRC to provide technical assistance in various aspects of computer technology that apply to computer-based reactor protection systems. This has involved the review of safety aspects of new reactor designs and the provision of technical advice on the use of computer technology in systems important to reactor safety. The latter includes determining and documenting state-of-the-art subjects that require regulatory involvement by the NRC because of their importance in the development and implementation of digital computer safety systems. These subjects include data communications, formal methods, testing, software hazards analysis, verification and validation, computer security, performance, software complexity and others. One topic — software reliability and safety — is the subject of this paper.

2. Methodology

The purpose of the software reliability task has been for LLNL to assist the NRC in understanding the state of the art in assessing the reliability of the software for a computer-based reactor protection system. Three separate activities were carried out: assess the status of national and international standards that relate

to software reliability and safety, obtain advice from technical experts in software reliability and safety, and assess the best current practice in industry. The first activity was reported in [2], so is not discussed further here.

2.1. Technical Experts

In 1992, LLNL sponsored a small workshop in which four leading software experts were brought together with members of the NRC staff to give their opinions on methods and techniques for the development of safety critical software and to discuss the reliability and safety of reactor software. Experts were chosen based on reputation, technical expertise, availability, and interest in the NRC's problem to achieve the shared goal of obtaining a variety of viewpoints. They were:

- Ricky Butler, NASA Langley.
- Nancy Leveson, University of California at Irvine¹.
- Bev Littlewood, City University of London.
- John Rushby, Stanford Research Institute.

The workshop was held in July 1992 in San Diego, California, and lasted a day and a half. A report on the workshop is available from LLNL [3].

2.2. Commercial Practice

In 1993, the investigation was expanded to study the best contemporary industry practices. Specific organizations within three companies were chosen based on reputation, availability, and willingness to participate. They were:

- Computer Sciences Corporation, Systems Engineering Division.
- International Business Machines, Federal Systems Company, Space Shuttle Project.
- TRW, Ballistic Missile Defense Division; Army WWMCCS Information System Project; and Universal Network Architecture Services Project. (Three projects within TRW were included in the study.)

Supplementary conversations were held with representatives of the NASA Software Engineering Laboratory and the American Institute of Aeronautics and Astronautics (AIAA) Software Reliability Project.

Interviews and discussions with the three companies were directed at "what works." Specifically, the discussions started with the following questions and expanded from there:

1. What are the most important obstacles to producing highly reliable software for use in safety-related applications?
2. Why is your company successful at producing highly reliable software?

¹ Now at the University of Washington.

3. What evidence exists that the methods work?
4. What are the most important lessons your company has learned about producing highly reliable software?

The results of the interviews and discussions were combined into a set of principles which were termed "design factors." This list was then discussed with a fourth company as a validation check. This company, which prefers to remain anonymous, is also highly regarded for the excellence of its software, but is representative of a very different segment of industry from that of the three primary companies. The results of this investigation have been documented and are available from LLNL [4].

3. Common Principles

Although the emphasis among the three sources of information (standards, experts and companies) tends to be quite different, there were no substantial areas of disagreement. This provides considerable confidence in the results of the study as applied to reactor protection systems and other safety-critical applications. Some broad, shared principles are given in this section. More detailed recommendations can be found in Section 4.

3.1. Safety Is a System Problem

This principle is generally accepted, but the implications for designers and regulators deserve some comment. Reactor protection systems are designed using the principles of diversity and defense-in-depth. The first implies that several different physical properties of the reactor are monitored for possible deviations from normality. Different logical combinations of sensed parameters are used to indicate problems, and different physical methods of response are used to prevent or mitigate failures. Thus, during an event or accident, a number of diverse signals that follow diverse paths through the equipment, will cause the reactor to trip, core cooling to operate, or containment isolation.

The second principle means that in typical reactor I&C designs there are multiple echelons of defense: the control and monitoring system is designed to keep the reactor out of trouble in the first place; the reactor trip system is designed to shut the reactor down when the operating envelope is exceeded; the safety parameter display system and manual controls are designed to provide enough information and capability for operator response; and the engineered safety features actuation system will actuate the emergency core cooling and containment isolation and cooling systems to cool the core and maintain containment integrity when required. At the physical barrier level, the core cladding is designed to prevent the release of radioactive material to the coolant, the pressure vessel and reactor coolant pressure boundary are designed to prevent the release of the coolant, and the containment is designed to prevent the release of radioactive material if the reactor coolant pressure boundary fails.

This discussion has some serious implications for software design. A balance should exist between the general architecture of the protection system and the software contained within it, so that functions are carried out at the most effective location within that architecture. For example, it is required by law that reactor protection systems have no single point of failure, so that multiple, redundant computer processors are required which must be resistant to simultaneous failure. Since computer systems can fail either because of hardware or software failures, a serious new concern (common-mode software failure) now exists where none existed before. A system-level response to this concern is using the system architecture to provide defense-in-depth (instead of attempting some form of n-version programming). An alternative is to ensure sufficient diversity among programs so that several programs do not execute almost-identical sequences of instructions on identical inputs, since similar faults in the programs can lead to common-mode failures.

3.2. Risk-Based Design

No company has unlimited resources — money, time, and (perhaps most important) highly qualified talent are all finite. It is also generally true that some portions of a software system have higher risk (in terms of the consequences of failure) than others. Indeed, it is frequently good practice to design software so that those portions whose failure can lead to the most adverse consequences are small, isolated and confined. Such software can be termed a “safety kernel.”

Risk-based design means that a formal risk analysis [5] is performed to identify the most risky portions of the software, to design both the overall protection system and its software to isolate and confine risk, and to concentrate resources on the areas of highest risk.

3.3. An Understanding of Complexity

The term “complexity” can be applied to many aspects of the software development process and products, and some of these are naturally more important than others. For the purposes of this paper, three types of complexity are recognized: functional complexity, structural (or design) complexity, and code complexity, with the caveat that all three do not have absolute definitions.

Functional complexity refers to the difficulty of the problem being solved — the functions that the software is expected to provide, and the interactions among those functions. This is very difficult to measure, so may receive insufficient attention.

Structural (or design) complexity refers to the complexity of the software architecture and design. The design must of course be sufficient to cover all the functional requirements, but should not be more complex than necessary.

Code complexity refers to the internal complexity of the code modules. This is the easiest to measure, so generally receives the majority of attention.

The fundamental issue is that complexity becomes a problem when it decreases the understandability or the verifiability of the computer system. No complexity aspect can be ignored, since concentrating on a single aspect just encourages designers to shift complexity elsewhere. For example, demanding that code modules have cyclomatic complexity metric less than ten is likely to result in an overly complex structure of modules, each of which has cyclomatic complexity metric less than ten.

3.4. Commitment to Quality

Perhaps the most outstanding shared characteristic of the companies interviewed was an understanding of, and a top-to-bottom commitment to, the production of quality software. All three companies demonstrated in various ways their pride in the quality of the software they build, and reward their people accordingly.

Quality did not come easily or painlessly. Current levels of achievement have taken a decade or more to reach. Mistakes have been made, but the attitude is "Why did we make the mistake?" and "How can we avoid it in the future?" rather than "Who shall we blame this time?"

Finally, none of the three companies is satisfied with its current status. Instead, continuous deliberate planned quality improvement appears to have become an integral part of their corporate cultures.

3.5. Multi-Level Evaluation

It is well accepted, at least among software practitioners, that there is no "silver bullet" [6] that can be used to magically produce reliable software. There is likewise no silver bullet for regulators to use in satisfying themselves that a software-based protection system is safe.

One of the principal results from the work at LLNL is the belief that assessing the safety of any software system involves at least three levels of evaluation: the development company and its culture, the specific development process for the product under review, and the software and documentary products of the development process.

The effect is to attempt to assure safety through the cumulative weight of diverse forms of evidence obtained primarily during the development process. This provides a qualitative argument for safety, as contrasted to a quantitative statistical assessment of the safety of the software system, such as the probability of failures per demand (f/d). This appears to be the best that can be done with the current state of software technology. There is a further implication that no guarantee can be given of absolute safety — but that is true of all technologies.

Since much of the evidence available is qualitative in nature, it may be quite difficult to calculate an accurate measure of safety, although crude bounds may

be possible. That is, it may not be possible to calculate the probability of an accident due to software failure to within a factor of two, but it may be possible to show that the probability is less than (say) 10^{-3} or 10^{-4} f/d, and add a well-founded belief through the weight of evidence that the probability is *much* lower than 10^{-3} f/d. This may be satisfactory if backed up by "an intellectually convincing argument."²

4. Design Factors

Given the failure of magic, the lack of "silver bullets," and the limitations of quantitative reasoning, confidence in the quality of any software product can only be acquired through years of proven successful use or a preponderance of a variety of evidence. The first alternative is unattractive (and most likely unattainable) when safety is important and is unlikely to exist in the rapidly changing computer industry. This leaves the challenging alternative of evaluating multiple factors relating to the software development company, design process, and software product. Success in other assessment arenas, such as TickIT [7] in the United Kingdom, provides additional confidence in this approach.

The research reported here has identified a number of design factors that can serve as the basis for a safety assessment. Supporting evidence comes from standards, expert opinion as given in the workshop and the technical literature, and the best current industry practice. Positive and negative design factors have been identified. They are listed here under five headings: general, management, process, product, and negative factors.

4.1. General Design Factors

The factors listed here relate generally to the development of safety-critical software.

- All levels of the company demonstrate a commitment to quality.
- There is longevity in personnel, policy and process which provides stability in company culture and continuous improvement over many years.
- Configuration management is implemented rigorously and used extensively during all life cycle phases and for all life cycle products, including test, verification and validation (V&V) and quality assurance (QA) products.
- Testing, verification and validation, and quality assurance activities are independent from the development activities.

² Our thanks to John Rushby for this phrase.

- The test team is involved from the very beginning of software development, to assure that the requirements are testable, to assure that the software to be developed employs design features that are within the capability of the test tools, and to provide input to the test cases, including test data and operational profiles.
- The test team and the V&V team include subject area specialists as well as software experts.
- Reviews, walkthroughs, and inspections are used at all stages of development and for all products, including testing, V&V, and QA products.
- Continuous process improvement is a corporate way of life.
- The company has long-term experience developing safety-critical applications.
- Software is one of the main businesses of the company, and top management understands the process of software development.
- The company has a goal of defect-free software, and spends the energy and resources to come as close to realizing the goal as humanly possible.
- Quality is built in. It is not possible to “test in” quality. Instead, quality must be designed into the product and that fact should be demonstrated by testing, V&V, and QA activities.

4.2. Management Design Factors

The factors listed here relate to company management of software development efforts. The underlying principle is an intelligent, deliberate commitment to quality.

4.2.1. General Management Design Factors

- Vendors, products, and services obtained from others are certified to the level required to support safety-critical applications.
- The company monitors, understands, and adapts to the changing environment in the computer industry.

4.2.2. Process Control Design Factors

- The company has a well-defined detailed software development process model.

- The software development process is stabilized through measurement, feedback, and gradual improvement.
- Process improvement occurs in two stages: a stabilization stage to provide a firm foundation for measurement of the effects of one (or at most a few) changes, which in turn is followed by measurement of the stabilized process outputs.
- Process data which is appropriate to the maturity of the company is collected, understood, and used as the basis for process improvement.

4.2.3. Management Action Design Factors

- The company's reward structure matches its commitment to quality, rather than undercutting it by primarily rewarding the meeting of budget and schedule targets.
- The company management actually uses its process models.
- Management predicts the effect of process changes through its process model, and measures results to obtain feedback on the actual effect of changes. In the language of control theory, the measurement lag would make an extremely sluggish system or an unstable one without the anticipation provided by predictions.
- Management has a successful track record planning, allocating resources, and meeting schedules within cost and quality constraints.
- Management actively identifies, analyzes, and manages business, technical and safety risks.
- Management abandons methods that don't work. (This may seem obvious, but abandoning an inappropriate "improvement" can be a career risk in some organizations.)
- Management ensures the planning, production and control of documentation. Because documentation tends to be neglected under stress, this serves as a sensitive indicator of management commitment.
- Management invites external review.
- Management understands that improvement takes time — typically, it takes about two years for a process change to become stabilized. Impatience and the search for short-term gains is not a virtue.

4.2.4. Personnel Design Factors

- Programming skill is not enough when safety is critical — some people must be skilled in the problem domain.

- It is well recognized that the single greatest factor in assuring quality software is the knowledge, skill, and intellectual ability of the technical staff.
- The greatest single obstacle to producing high-quality software is inaccurate interpersonal communications. As a corollary, project teams should be kept small (6–8 people).

4.3. Process Design Factors

The factors listed here relate to both management and technical aspects of the software development process. The underlying principle is a deliberate planned approach to software development.

4.3.1. General Process Design Factors

- Reviews, walkthroughs, and inspections are used at all stages of development and for all products, including code, development documents, verification and validation (V&V) products and quality assurance (QA) products.
- Resource investments are targeted throughout the process life cycle according to safety and reliability requirements.

4.3.2. Technical Planning Design Factors

- An appropriate life cycle model is deliberately chosen and used.
- Life cycle activities are chosen to promote early detection of errors.
- The software architecture is chosen to isolate and confine risks, and the software elements are managed appropriately according to risk.

4.3.3. Requirements Specification Design Factors

- Requirements are stable. For a reactor protection system, unstable requirements reflect inadequate system design.
- Requirements are analyzed to understand their implications: to detect inconsistencies, unneeded but expensive specifications, over-specification, and requirements that may be extremely difficult or impossible to fulfill; and to ensure that the requirements are correctly translated from application-specific terminology to software-specific terminology.
- Requirements are validated against the system design and the safety analysis report.

- Prototyping and simulation are used in appropriate ways to refine requirements, test design approaches and demonstrate system performance.
- Requirements are written to be testable.
- The requirements analyst prepares and documents a well-constructed argument as to why the requirements are correct and complete.

4.3.4. Design Specification Design Factors

- Safety-critical software components are identified early in the design process so that sufficient resources can be directed to them.
- A design philosophy suitable for safety-critical software is used. In particular, "risky" practices are avoided.
- An appropriate level of complexity is defined for the product and documented practices are followed that control this complexity.
- The product is designed to permit easy understanding, testing, and verification.
- The designer presents and documents a well-constructed argument as to why the design is correct, complete, reliable, and safe.

4.3.5. Software Quality Assurance Design Factors

- The software quality assurance effort is organizationally independent of the development organization.
- Quality is built in. It cannot easily be retrofitted, and cannot be "tested in."
- Defect tracking is taken seriously, is carried out uniformly and consistently, and is statistically valid.
- The root causes of defects are determined, and appropriate corrective actions are initiated to reduce the probability of similar errors in the future.

4.3.6. Safety Design Factors

- Hazards analysis is used as part of the development process. Hazards can be introduced by the selection of design approaches, certain hardware, software tools, or the use of software itself as a solution to a safety problem.
- System diversity is used to improve reliability and safety, and the software is designed to be an appropriate part of the system.

- The safety system of which software is a part may be circumvented, turned off, or driven into failure by operator actions. Neither the system nor the software can be expected to prevent these problems. These implications are factored into the software design.
- Reliability in the order of 10^{-5} to 10^{-6} failures per demand ³ cannot be assured by any known method of implementation or testing. The software portion of the system must either be designed in such a way that these levels of reliability are not required, or the inability to quantify reliability at these levels must be accepted.

4.3.7. Testing Design Factors

- Testing is carried out at multiple levels: unit, subsystem, system.
- The software testing effort is organizationally independent of the development organization.
- Testing has its own life cycle, which is planned, designed, implemented, and followed in parallel to the development life cycle.
- Testing goals do not exceed the current practical limitation of about 10^{-3} to 10^{-4} failures per demand.

4.3.8. Verification and Validation Design Factors

- The software verification and validation effort is organizationally independent of the development organization.
- Requirements validation is performed.
- V&V is planned early in the life cycle, and the V&V plan is peer-reviewed.
- The software product is designed to permit effective verification and validation.

4.4. Product Design Factors

The half dozen factors listed here are aimed at creating a predictable, verifiable software system. The underlying principles are simplicity and determinism. The presence of these factors is considered a positive indication of lowered complexity or easier error detection.

- *No interrupts.* The use of interrupts, beyond a simple clock interrupt, is considered a higher-risk implementation method because of the extra

³ Roughly equivalent to 10^{-8} to 10^{-9} failures per year in continuously operating vehicles such as aircraft.

care required to ensure correct synchronization between interrupt code and interrupted code, and to ensure that interrupted code is correctly resumed.

- *No multi-tasking.* Multi-tasking requires context switching and task management in addition to the complications attendant upon using interrupts.
- *Simple loop design.* A single loop program structure is the simplest program organization capable of continuous operation.
- *Deterministic predictable timing.* Evidence should exist that software product timing is a predictable function of load, and that load is limited by design.
- *No pointers.* The use of explicit pointers (addresses) of data has been taken by some as a risky practice. The potential exists for errors in programmer-directed address arithmetic which would not exist if named variables were used and the addresses were computed automatically by compiler.
- *Strong data typing.* Data typing permits compilers to detect data misuse errors (e.g., using an integer as if it were a floating point number). This class of error represents a significant proportion of all errors made, and strong data typing with good compilers almost eliminates it.

4.5. Negative Factors

The presence of any of these factors provides ample justification for caution and concern by the regulator. A more thorough investigation may well be in order.

- *High turnover.* The most obvious implication of high turnover is that building a team of high-quality people with a team memory is impossible. Less obvious is the fact that high turnover is a comment (by software engineers and managers who leave) on the competence of management that is left behind. It should not be ignored.
- *Projects are schedule-driven.* The first victims of a missed deadline are usually quality assurance and documentation. The next victim is the testing program. A "deliver at all costs" mentality is cause for concern.
- *Organizational process history is short or lacking.* The research results are explicit about the length of time it takes to build a quality software organization.
- *Management cannot enforce stable requirements.* Stable and complete requirements are necessary for quality software products, but the role of management in ensuring this cannot be emphasized enough. Not only must management demand that requirements be locked down, but

management itself must not be the source of requirements thrashing. Requirements instability and weak management control are indicators of potential failures.

- *Management has a record of failing to meet predicted cost, schedule, and quality goals for products.* This is typically an indication of management by chaos or paradoxically, schedule-driven rather than quality-driven development. Schedule- and budget-driven development schemes often fail to meet delivery schedules because of product non-performance problems. Something is delivered, but it is not the contracted-for product. A record of failing to meet cost, schedule, and quality goals should be taken seriously as an indicator of deeper troubles.
- *The organization fails to track errors and defect causes.* An organization's record of errors, causes, and corrective actions is its won-loss track record. No record, or a haphazard record, should be taken in default as meaning a bad record.
- *The effort is underfunded.* Several of the companies interviewed suggested that most large government software contracts are underfunded by at least a factor of two with the expectation that more funds can be obtained later by litigation, contract expansion, or cost overrun procedures. Whatever the reason, underfunding results in staff transients and failures to carry out "non-essential" but vital activities such as quality assurance, documentation, and V&V. While it may be difficult for an outside reviewer to estimate what a correct funding profile should be, this negative factor is very real.
- *The organization exhibits "kill the messenger" syndrome.* Several companies interviewed emphasized the need for administrative procedures by which bearers of bad tidings could unburden themselves without jeopardizing their careers. They noted that organizations without these mechanisms were often the last to know about internal problems.

5. Conclusions

It is certainly true that unsafe software has been created, and will probably be created in the future. It is equally true that very good software has been created, and will continue to be created in the future. As with any technology controlling a safety-critical application, software has contributed to death and injury [8], but it is somewhat surprising how rare this has been. Why?

The thesis presented in this paper is that much of the difference between success and failure can be attributed to the knowledge, understanding, intelligence and care of the individuals and companies involved in the development of safety-critical software. By combining the best from theory and practice it is possible to isolate a number of factors that distinguish the good from the bad.

There is much emphasis in the literature, and in conferences such as this one, on the negative consequences of software failures in safety-critical applications and this is appropriate. However, as suggested in this paper, there are also several balancing, positive factors which deserve equal emphasis. In particular, software does not wear out, can potentially be used to identify and compensate for hardware failures, and can potentially provide much greater control to operators during unexpected events. These factors should be carefully evaluated on a case-by-case basis to determine the suitability of software in each plant application.

The challenge faced by software developers is to use software safely to increase the reliability of the application, while the challenge for the NRC evaluator is to assure that this is done. The research conducted as part of this contractual effort suggests that convincing evidence can be obtained in practice that reliable safety-critical software is being or has been developed. However, neither the development of such software nor the effort required to certify it for safety-critical usage is easy.

References

1. "On a plateau after 50 years," *Nuclear Engineering International* **38**, #467 (June 1993), 18–24.
2. Debra Sparkman, "Standards and practices for reliable safety-related software systems," *Proc. Third International Symposium on Software Reliability Engineering*, IEEE Computer Society (October 1993), 318–328.
3. J. Dennis Lawrence, "Workshop on developing safe software: final report," UCRL-ID-113438, Lawrence Livermore National Laboratory, Livermore, CA (November 30, 1992).
4. J. Dennis Lawrence and Warren L. Persons, "A survey of industry practices in developing high-integrity software", Lawrence Livermore National Laboratory, Livermore, CA, in prep.
5. Robert N. Charette, *Software Engineering Risk Analysis and Management*, McGraw-Hill (1989).
6. Frederick P. Brooks, "No silver bullet: essence and accidents of software engineering," *IEEE Computer* **20**, 4 (April 1987), 10–19.
7. "Guide to software quality management system construction and certification," TickIT Project Office, London (February 28, 1992).
8. Nancy G. Leveson and Clark S. Turner, "An investigation of the Therac-25 accidents," *IEEE Computer* **26**, 7 (July 1993), 18–41.

**DATE
FILMED**

6/17/94

END