Analyzing Software Requirements Errors in Safety-Critical, Embedded Systems

Robyn R. Lutz*
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109

Abstract

This paper analyzes the root causes of safety-related software errors in safety-critical, embedded systems. The results show that software errors identified as potentially hazardous to the system tend to be produced by different error mechanisms than non-safety-related software errors. Safety-related software errors are shown to arise most commonly from (1) discrepancies between the documented requirements specifications and the requirements needed for correct functioning of the system and (2) misunderstandings of the software's interface with the rest of the system. The paper uses these results to identify methods by which requirements errors can be prevented. The goal is to reduce safety-related software errors and to enhance the safety of complex, embedded systems.

I. Introduction

This paper examines 387 software errors uncovered during integration and system testing of two spacecraft, Voyager and Galileo. A software error is defined to be a software-related discrepancy between a computed, observed, or measured value or condition and the true, specified, or theoretically correct value or condition [6]. Each of these software errors was documented at the time of discovery by a form describing the problem or failure. The form also recorded the subsequent analysis and the corrective actions taken.

As part of the standard procedure for correcting each reported software error, the failure effect of each is classified as negligible, significant, or catastrophic. Those classified as significant or catastrophic are investigated by a systems safety analyst as representing potential safety hazards [13]. For this study the 87 software errors on Voyager and 122 software errors on Galileo documented as having potentially significant or catastrophic effects are classified as safety-related.

The spacecrafts' software is safety-critical in that it monitors and controls components that can be involved in hazardous system behavior [11]. The software must execute in a system context without contributing unacceptable risk.

Each spacecraft involves embedded software distributed on several different flight computers. Voyager has roughly 18,000 lines of source code; Galileo has over 22,000 [18]. Embedded software is software that runs on a computer system that is integral to a larger system whose primary purpose is not computational [6]. The software on both spacecraft is highly interactive in terms of the degree of message-passing among system components, the need to respond in real-time to monitoring of the hardware and environment, and the complex timing issues among parts of the system. The software development teams for each spacecraft involved multiple teams working for a period of years.

The purpose of this paper is to identify the extent and ways in which the cause/effect relationships of safety-related software errors differ from the cause/effect relationships of non-safety-related software errors. In particular, the analysis shows that errors in identifying or understanding functional and interface requirements frequently lead to safety-related software errors. This distinction is used to identify methods by which the common causes of safety-related software errors can be targeted during development. The goal is to improve system safety by understanding and, where possible, removing, the prevalent sources of safety-related software errors.

---

*Author's mailing address is Dept. of Computer Science, Iowa State University, Ames, IA 50011. This paper appeared in RE '93, the *Proceedings of the IEEE International Symposium on Requirements Engineering*, Jan 4-6, 1993, San Diego, CA. The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA.

II. Methodology

The study described here characterizes the root causes of the safety-related software errors discovered during integration and system testing. The recent work by Nakajo and Kume on software error cause/effect relationships offers an appropriate framework for classifying the software errors [16]. Their work is extended here to account for the additional complexities operative in large, safety-critical, embedded systems with evolving requirements driven by hardware and environmental issues.

Nakajo and Kume's classification scheme analyzes three points in the path from a software error backwards to its sources. This approach allows classification not only of the documented software error (called the program fault), but also of the earlier human error (the root cause, e.g., a misunderstanding of an interface specification), and, before that, of the process flaws that contribute to the likelihood of the error's occurrence (e.g., inadequate communication between systems engineering and software development teams).

The classification scheme thus leads backwards in time from the evident software error to an analysis of the root cause (usually a communication error or an error in recognizing or deploying requirements), to an analysis of the software development process. By comparing common error mechanisms for the software errors identified as potentially hazardous with those of the other software errors, the prevalent root causes of the safety-related errors are isolated. The classification of the sources of error is then applied here to determine countermeasures which may prevent similar error occurrences in other safety-critical, embedded systems. This paper thus uses the classification scheme to assemble an error profile of safety-related software errors and to identify development methods by which these sources of error can be controlled.

III. Analysis of Safety-Related Software Errors

A. Overview of Classification Scheme

An overview of the classification scheme follows, adjusted to the needs of safety-critical, embedded software. See [16] for additional details on how errors are categorized. An ongoing, multi-project investigation will address the issue of repeatability (do different analysts classify a given error in the same way?).

- Program Faults (Documented Software Errors)

  A. Internal Faults (e.g., syntax)

  B. Interface Faults (interactions with other system components, such as transfer of data or control)

  C. Functional Faults (operating faults: omission or unnecessary operations; conditional faults: incorrect condition or limit values; behavioral faults: incorrect behavior, not conforming to requirements)

- Human Errors (Root Causes)

  A. Coding or Editing Errors

  B1. Communication Errors Within a Team (misunderstanding S/W interface specifications)

  B2. Communication Errors Between Teams (misunderstanding H/W interface specifications or other team's S/W specifications)

  C1. Errors in Recognizing Requirements (misunderstanding specifications or problem domain)

  C2. Errors in Deploying Requirements (problems implementing or translating requirements into a design)

- Process Flaws (Flaws in Control of System Complexity + Inadequacies in Communication or Development Methods)

  A. Inadequate Code Inspection and Testing Methods

  B1. Inadequate Interface Specifications + Inadequate Communication (among S/W developers)

  B2. Inadequate Interface Specifications + Inadequate Communication (between S/W and H/W developers)

  C1. Requirements Not Identified or Understood + Incomplete Documentation

  C2. Requirements Not Identified or Understood + Inadequate Design

Clearly, the attribution of a key human error and a key process flaw to each software error oversimplifies the cause/effect relationship. However, the identification of these factors allows the characterization of safety-related software errors in a way that relates features of the development process and of the system under development to the safety consequences of those features. Similarly the association of each software error with a human error, while unrealistic (in what sense is a failure to predict details of system behavior an error?), allows a useful association between human factors (such as misunderstanding the requirements or the underlying physical realities) and their safety-related consequences.

B. Program Faults

Safety-related software errors account for 56% of the total software errors for Voyager and 48% of the total software errors for Galileo discovered during integration and system testing. (The reader is referred to [14] for the tables containing the error data.)

Few *internal* faults (e.g., coding errors internal to a software module) were uncovered during integration and system testing. An examination of software errors found later during operations also shows few internal errors. It appears that these coding errors are being detected and corrected before system testing begins. They are thus not discussed further in this paper.

At a high level of detail, safety-related and non-safety-related software errors display similar proportions of interface and functional faults. *Functional faults* (operating, conditional, or behavioral discrepancies with the functional requirements) are the most common kind of software error. Behavioral faults account for about half of all the functional faults on both spacecraft (52% on Voyager; 47% on Galileo).

On Voyager fully half of the safety-related functional faults are attributable to behavioral faults (the software behaving incorrectly). On Galileo, a slightly greater percentage is due to operating faults (nearly always a required but omitted operation in the software) than to behavioral faults. Often the omitted operation involves the failure to perform adequate reasonableness checks on data input to a module. This frequently results in an error-recovery routine being called inappropriately.

Conditional faults (nearly always an erroneous value on a condition or limit) tend to be safety-related on both spacecraft (73% total). Erroneous values (e.g., of deadbands or delay timers) often involve risk to the spacecraft by causing inappropriate triggering of an error-recovery response or by failing to trigger a needed response. The association between conditional faults and safety-related software errors emphasizes the importance of specifying the correct values for any data used in control decisions in safety-critical, embedded software.

The analysis also identifies *interface faults* (incorrect interactions with other system components, such as the timing or transfer of data or control) as a significant problem (36% of the safety-related program faults on Voyager; 19% on Galileo). Sect. IV below describes how the high incidence of interface faults in these complex, embedded systems contrasts with the low incidence of interface faults in earlier studies on simpler, standalone software.

C. Relationships Between Program Faults and Root Causes

The second step in the cause/effect analysis is to trace backwards in time to the human factors involved in the program faults that were discovered during integration and system testing.

For *interface faults*, the major human factors are either communication errors within a development team or communication errors between a development team and other teams. In the latter case, a further distinction is made between misunderstanding hardware interface specifications and misunderstanding the interface specifications of other software components. Communication errors between development teams (rather than within teams) is the leading cause of interface faults (93% on Voyager, 72% on Galileo). Safety-related interface faults are associated overwhelmingly with *communication errors* between a development team and others (often between software developers and systems engineers), rather than with communication errors within a team.

Significant differences in the distribution of fault causes between safety-related and non-safety-related interface faults appear. The primary cause of safety-related *interface* faults is *misunderstood hardware interface specifications* (67% on Voyager; 48% on Galileo). Examples are faults caused by wrong assumptions about the initial state of relays or by unexpected heartbeat timing patterns in a particular operating mode. (On the other hand, the root causes of non-safety-related interface faults are distributed more evenly between misunderstood hardware specifications and misunderstood software specifications.) The profile of safety-related interface faults assembled here emphasizes the importance of understanding the software as a set of embedded components in a larger system.

The primary cause of safety-related *functional* faults is errors in *recognizing (understanding) the requirements* (62% on Voyager, 79% on Galileo). Safety-related conditional faults, for example, are almost always caused by errors in recognizing requirements. (On the other hand, non-safety-related functional faults are more often caused by errors in deploying—implementing—the requirements.)

In summary, difficulties with requirements is the key root cause of the safety-related software errors which have persisted until integration and system testing.

D. Relationships Between Root Causes and Process Flaws

In tracing backwards from the program faults to their sources, features of the system-development process can be identified which facilitate or enable the occurrence of errors. Discrepancies between the difficulty of the problem and the means used to solve it may permit hazardous software errors to occur [4].

The third step of the error analysis therefore associates a pair of process flaws with each program fault [16]. The first element in the pair identifies a process flaw or inadequacy in the *control of the system complexity* (e.g., requirements which are not discovered until system testing). The second element of the pair identifies an associated process flaw in the *communication or development methods used* (e.g., imprecise or unsystematic specification methods). The two elements of the process-flaw pair are closely related. Frequently, as will be seen in Sect. V, a solution to one flaw will provide a solution to the related flaw.

For safety-related *interface faults*, the most common complexity-control flaw is *interfaces not adequately identified or understood* (56% on Voyager; 87% on Galileo). The most common safety-related flaw in the communication or development methods used on Voyager is *hardware behavior not documented* (44%). On Galileo the most common safety-related flaws are *lack of communication between H/W and S/W teams* (35%) and *interface specifications known but not documented or communicated* (35%).

Anomalous hardware behavior is a more significant factor in safety-related than in non-safety-related interface faults. It is often associated with interface design during system testing, another indication of a unstable software product.

There are significant variations in the process flaws that cause errors between the two spacecraft. Interface design during testing is involved in almost one-fifth of the safety-critical interface faults on Voyager, but in none of them on Galileo. This is because on Voyager a set of related hardware problems generated nearly half of the safety-related interface faults. On the other hand, the problem of interface specifications that are known but not documented is more common on Galileo. This is perhaps due to the increased complexity of the Galileo interfaces.

For *functional faults*, requirements not identified and requirements not understood are the most common complexity-control flaws. Safety-related functional faults are more likely than non-safety-related functional faults to be caused by *requirements which have not been identified.*

With regard to flaws in the communication or development methods, *missing requirements* are involved in nearly half the safety-related errors that involve recognizing requirements. *Inadequate design* is the most common flaw leading to errors in deploying requirements on Voyager. On Galileo, *incomplete documentation of requirements* is as important a factor for safety-related errors, but not for non-safety-related errors.

Imprecise or unsystematic specifications are more than twice as likely to be associated with safety-related functional faults as with non-safety-related functional faults. Similarly, *unknown, undocumented, or wrong requirements* are a greater cause of safety-related than of non-safety-related errors.

These results suggest that the sources of safety-related software errors lie farther back in the software development process–in inadequate requirements–whereas the sources of non-safety-related errors more commonly involve inadequacies in the design phase.

IV. Comparison of Results with Previous Work

Although software errors and their causes have been studied extensively, the current work differs from most of the prior investigations in the four following ways:
1) The software chosen to analyze in most studies is not embedded in a complex system as it is here. The consequence is that the role of interface specifications in controlling software hazards has been underestimated.
2) Unlike the current paper, most studies have analyzed fairly simple systems in familiar and well-understood application domains. Consequently, few software errors have occurred during system testing in most studies, leading to a gap in knowledge regarding the sources of these more-persistent and often more hazardous errors.
3) Most studies assume that the requirements specification is correct. On the spacecraft, as in many large, complex systems, the requirements evolve as knowledge of the system's behavior and the problem domain evolve. Similarly, most studies assume that requirements are fixed by the time that systems testing begins. This leads to a underestimation of the impact

of unknown requirements on the scope and schedule of the later stages of the software development process. 4) The distinction between causes of safety-critical and non-safety-critical software errors has not been adequately investigated. Efforts to enhance system safety by specifically targeting the causes of safety-related errors, as distinguished from the causes of all errors, can take advantage of the distinct error mechanisms, as described in Sect. 5.

A brief description of the scope and results of some related work is given below and compared with the results presented in this paper for safety-critical, embedded computer systems.

Nakajo and Kume categorized 670 errors found during the software development of two firmware products for controlling measuring instruments and two software products for instrument measurement programs [16]. Over 90% of the errors were either interface or functional faults, similar to the results reported here. While the key human error on the spacecraft involved communication between teams, the key human error in their study involved communication within a development team. Finally, the key process flaw that they identified was a lack of methods to record known interfaces and describe known functions. In the safety-critical, embedded software on the spacecraft, the flaw was more often a failure to identify or to understand the requirements.

Ostrand and Weyuker categorized 173 errors found during the development and testing of an editor system [19]. Only 2% of the errors were found during system testing, reflecting the simplicity and stability of the interfaces and requirements. Most of the errors (61%) were found instead during function testing. Over half of these errors were caused by omissions, confirming the findings of the present study that omissions are a major cause of software errors.

Schneidewind and Hoffmann [21] categorized 173 errors found during the development of four small programs by a single programmer. Again, there were no significant interfaces with hardware and little system testing. The most frequent class of errors, other than coding and clerical, was design errors. All three of the most common design errors–extreme conditions neglected, forgotten cases or steps, and loop control errors– are also common functional faults on the spacecraft.

Both the findings presented in [19, 21] and in this paper confirm the common experience that early insertion and late discovery of software errors maximizes the time and effort that the correction takes. Errors inserted in the requirements and design phases take longer to find and correct than those inserted in later phases (because they tend to involve complex software structures). Errors discovered in the testing phase take longer to correct (because they tend to be more complicated and difficult to isolate). This is consistent with the results in [18] indicating that more severe errors take longer to discover than less severe errors during system-level testing. Furthermore, this effect was found to be more pronounced in more complex (as measured by lines of code) software.

The work done by Endres is a direct forerunner of Nakajo and Kume's in that Endres backtracked from the error type to the technical and organizational causes which led to each type of error [4]. Moreover, because he studied the system testing of an operating system, the software's interaction with the hardware was a source of concern. Endres noted the difficulty of precisely specifying functional demands on the systems before the programmer had seen their effect on the dynamic behavior of the system. His conclusion that better tools were needed to attack this problem still holds true eighteen years after he published his study.

Of the 432 errors that Endres analyzed, 46% were errors in understanding or communicating the problem, or in the choice of a solution, 38% were errors in implementing a solution, and the remaining 16% were coding errors. These results are consistent with the finding here that software with many system interfaces displays a higher percentage of software errors involving understanding requirements or the system implications of alternative solutions.

Eckhardt et al., in a study of software redundancy, analyzed the errors in twenty independent versions of a software component of an inertial navigation system [3]. He found that inadequate understanding of the specifications or the underlying coordinate system was a major contributor to the program faults causing coincident failures.

Addy, looking at the types of errors that caused safety problems in a large, real-time control system, concluded that the design complexity inherent in such a system requires hidden interfaces which allow errors in non-critical software to affect safety-critical software [1]. This is consistent with Selby and Basili's results when they analyzed 770 software errors during the updating of a library tool [22]. Of the 46 errors documented in trouble reports, 70% were categorized as "wrong" and 28% as "missing." They found that subsystems that were highly interactive with other subsystems had proportionately more errors than less interactive subsystems.

Leveson listed a set of common assumptions that are often false for control systems, resulting in software errors [11]. Among these assumptions are that the software specification is correct, that it is possible to predict realistically the software's execution environment (e.g., the existence of transients), and that it is possible to anticipate and specify correctly the software's behavior under all possible circumstances. These assumptions tend to be true for the simple systems in which software errors have been analyzed to date and false for spacecraft and other large, safety-critical, embedded systems. Thus, while studies of software errors in simple systems can assist in understanding internal errors or some functional errors, they are of less help in understanding the causes of safety-related software errors, which tend heavily to involve interfaces or recognition of complex requirements.

Similarly, standard measures of the internal complexity of modules have limited usefulness in anticipating software errors during system testing It is not the internal complexity of a module but the complexity of the module's connection to its environment that yields the persistent, safety-related errors seen in the embedded systems here [8].

V. Conclusion

A. Recommendations

The results in Sect. III indicate that safety-related software errors tend to be produced by different error mechanisms than non-safety-related software errors. This means that system safety can be directly enhanced by targeting the causes of safety-related errors. Specifically, the following six recommendations emerge from our analysis of safety-related errors in complex, embedded systems.

1. Focus on the interfaces between the software and the system in analyzing the problem domain, since these interfaces are a major source of safety-related software errors.

The traditional goal of the requirements analysis phase is the specification of the software's external interface to the user. This definition is inadequate when the software is deeply embedded in larger systems such as spacecraft, advanced aircraft, air-traffic control units, or manufacturing process-control facilities. In such systems, the software is often physically and logically distributed among various hardware components of the system.

Specifying the external behavior of the software (its transformation of software inputs into software outputs) only makes sense if the interfaces between the system inputs (e.g., environmental conditions, power transients) and the software inputs (e.g., monitor data) are also specified. Similarly, specifying the interfaces–especially the timing and dependency relationships–between the software outputs (e.g., star identification) and system outputs (e.g., closing the shutter on the star scanner) is necessary. [5, 10]

System-development issues such as timing (real-time activities, interrupt handling, frequency of sensor data), hardware capabilities and limitations (storage capacity, power transients, noise characteristics), communication links (buffer and interface formats), and the expected operating environment (temperature, pressure, radiation) need to be reflected in the software requirements specifications because they are frequently sources of safety-critical software interface errors.

Timing is a particularly difficult source of safety-related software interface faults since timing issues are so often integral to the functional correctness of safety-critical, embedded systems. Timing dependencies (e.g., how long input data is valid for making control decisions) should be included in the software interface specifications. Analytical models or simulations to understand system interfaces are particularly useful for complex, embedded systems.

2. Identify safety-critical hazards early in the requirements analysis.

These hazards are constraints on the possible designs and factors in any contemplated tradeoffs between safety (which tends to encourage software simplicity) and increased functionality (which tends to encourage software complexity) [10, 22]. Many of the safety-related software errors reported in Sect. III involve data objects or processes that would be targeted for special attention using hazard-detection techniques such as those described in [7, 11]. Early detection of these safety-critical objects and increased attention to software operations involving them might forestall safety-related software errors involving them.

3. Use formal specification techniques in addition to natural-language software requirements specifications.

Lack of precision and incomplete requirements led to many of the safety-related software errors seen here. Enough detail is needed to cover all circumstances that can be envisioned (component failures, timing constraint violations, expired data) as well as to document all environmental assumptions (e.g., how close to the sun an instrument will point) and assumptions about other parts of the system (maximum transfer rate, consequences of race conditions or cycle slippage). The capability to describe dynamic events, the timing of process interactions in distinct computers, decentralized supervisory functions, etc., should be considered in chooosing a formal method [2, 4, 5, 15, 20, 23].

4. Promote informal communication among teams.

Many safety-related software errors resulted from one individual or team misunderstanding a requirement or not knowing a fact about the system that member(s) of another development team knew. The goal is to be able to modularize responsibility in a development project without modularizing communication about the system under development. The identification and tracking of safety hazards in a system, for example, is clearly best done across team boundaries.

5. As requirements evolve, communicate the changes to the development and test teams.

This is both more important (because there are more requirements changes during design and testing) and more difficult (because of the number and size of the teams and the length of the development process) in a large, embedded system than in simpler systems. In analyzing the safety-related software errors, it is evident that the determination as to who needs to know about a change is often made incorrectly. Frequently, changes that appear to involve only one team or system component end up affecting other teams or components at some later date (sometimes as the result of incompatible changes in distinct units).

There is also a need for faster distribution of changes that have been made, with the update stored so as to be fingertip accessible. CASE tools offer a possible solution to the difficulty of promulgating change without increasing paperwork.

The prevalence of safety-related software errors involving misunderstood or missing requirements points up the inadequacy of consistency checks of requirements and code as a means of demonstrating system correctness [10]. Code that implements incorrect requirements is incorrect if it fails to provide needed system behavior.

Similarly, generating test cases from misunderstood or missing requirements will not test system correctness. Traceability of requirements and automatic test generation from specifications offers only partial validation of complex, embedded systems. Alternative validation and testing methods such as those described in [9, 11] offer greater coverage.

6. Include requirements for "defensive design" [17].

Many of the safety-related software errors involve inadequate software responses to extreme conditions or extreme values. Anomalous hardware behavior, unanticipated states, events out of order, and obsolete data are all causes of safety-related software errors on the spacecraft.

Run-time safety checks on the validity of input data, watchdog timers, delay timers, software filters, software-imposed initialization conditions, additional exception handling, and assertion checking can be used to combat the many safety-critical software errors involving conditional and omission faults [11]. Requirements for error-handling, overflow protection, signal saturation limits, heartbeat and pulse frequency, maximum event duration, and system behavior under unexpected conditions can be added and traced into the design. Many safety-related functional faults involve error-recovery routines being invoked inappropriately because of erroneous limit values or bad data.

Backward analysis from critical failures to possible causes offers one check of how defensive the requirements and design are [12]. Requirements specifications that account for worst-case scenarios, models that can predict the range of possible (rather than allowable) values, and simulations that can discover unexpected interactions before system testing contribute to the system's defense against hazards.

B. Summary and Future Work

In large, embedded systems such as the two spacecraft in this study, the software requirements change throughout the software development process, even during system testing. This is largely due to unanticipated behavior, dynamic changes in the operating environment, and complex software/hardware and software/software interactions in the systems being developed. Controlling requirement changes (and, hence, the scope and cost of development) is difficult since the changes are often prompted by an improved understanding of the software's necessary interfaces with the physical components of the spacecraft in which it is embedded. Complex timing issues and hardware idiosyncrasies often prompt changes to requirements or to design solutions.

The analysis presented here of the cause/effect relationships of safety-related software errors pinpoints aspects of system complexity which merit additional attention. Specifically, the results have shown that conditional faults (e.g., condition or limit values) are highly correlated with safety-related software errors. Operating faults (especially the omission of run-time reasonableness checks on data) are also highly correlated with safety-related software errors. Unknown, undocumented, or erroneous requirements frequently are associated with safety-related software errors as well. Hardware/software interfaces have been shown to be a frequent trouble spot because of the lack of communication between teams.

The results presented in this paper indicate a need for better methods to confront the real-world issues of developing safety-critical, embedded software in a complex, distributed system. Future work will be directed at incorporating knowledge of the distinct error mechanisms that produce safety-related software errors into the requirements analysis and validation processes. Work is also needed on specifying how these results can be used to predict more precisely what features or combinations of factors in a safety-critical, embedded system are likely to cause time-consuming and hazardous software errors.

References

[1] E. A. Addy, "A Case Study on Isolation of Safety-Critical Software," in *Proc 6th Annual Conf on Computer Assurance*. NIST/IEEE, 1991, pp. 75–83.

[2] A. M. Davis, *Software Requirements, Analysis and Specification.* Englewood Cliffs, N.J.: Prentice Hall, 1990.

[3] D. E. Eckhardt, et al., "An Experimental Evaluation of Software Redundancy as a Strategy for Improving Reliability," *IEEE Trans Software Eng*, 17, 7, July 1991, pp. 692–702.

[4] A. Endres, "An Analysis of Errors and Their Causes in Systems Programs," *IEEE Trans Software Eng*, SE-1, 2, June 1975, pp. 140–149.

[5] E. M. Gray and R. H. Thayer, "Requirements," in *Aerospace Software Engineering, A Collection of Concepts.* Ed. C. Anderson and M. Dorfman. Washington: AIAA, 1991, pp. 89–121.

[6] ANSI/IEEE Standard Glossary of Software Engineering Terminology. New York: IEEE, 1983.

[7] M. S. Jaffe et al., "Software Requirements Analysis for Real-Time Process-Control Systems," *IEEE Trans Software Eng*, 17, 3, March 1991, pp. 241–258.

[8] P. Jalote, *An Integrated Approach to Software Engineering.* New York: Springer-Verlag, 1991.

[9] J. C. Knight, "Testing," in *Aerospace Software Engineering, A Collection of Concepts.* Ed. C. Anderson and M. Dorfman. Washington: AIAA, 1991, pp. 135–159.

[10] N. G. Leveson, "Safety," in *Aerospace Software Engineering, A Collection of Concepts.* Ed. C. Anderson and M. Dorfman. Washington: AIAA, 1991, pp. 319–336.

[11] N. G. Leveson, "Software Safety in Embedded Computer Systems," *Commun ACM*, Vol. 34, No. 2, Feb 1991, pp. 35–46.

[12] N. G. Leveson and P. R. Harvey, "Analyzing Software Safety," *IEEE Transactions on Software Engineering*, SE-9, 5, Sept 1983, pp. 569–579.

[13] Karan L'Heureux, "Software Systems Safety Program RTOP, Phase A Report," Internal Document, Jet Propulsion Laboratory, April 19, 1991.

[14] R. Lutz, "Analyzing Software Requirements Errors in Safety-Critical, Embedded Systems," TR92-27, Dept. of Comp. Sci, Iowa State University.

[15] R. Lutz and J. S. K. Wong, "Detecting Unsafe Error Recovery Schedules," *IEEE Trans Software Eng*, 18, 8, Aug, 1992, pp. 749–760.

[16] T. Nakajo and H. Kume, "A Case History Analysis of Software Error Cause–Effect Relationships," *IEEE Trans Software Eng* 17, 8, Aug 1991, pp. 830-838.

[17] P. G. Neumann, "The Computer-Related Risk of the Year: Weak Links and Correlated Events," in *Proc 6th Annual Conf on Computer Assurance*. NIST/IEEE, 1991, pp. 5–8.

[18] A. P. Nikora, "Error Discovery Rate by Severity Category and Time to Repair Software Failures for Three JPL Flight Projects," Internal Document, Jet Propulsion Laboratory, 1991.

[19] T. J. Ostrand and E. J. Weyuker, "Collecting and Categorizing Software Error Data in an Industrial Environment," *The Journal of Systems and Software*, 4, 1984, pp. 289–300.

[20] *Proc Berkeley Workshop on Temporal and Real-Time Specification*. Eds. P. B. Ladkin and F. H. Vogt. Berkeley, CA: International Computer Science Institute, 1990, TR-90-060.

[21] N. F. Schneidewind and H.-M. Hoffmann, "An Experiment in Software Error Data Collection and Analysis," *IEEE Trans Software Eng*, SE-5, 3, May 1979, pp. 276–286.

[22] R. W. Selby and V. R. Basili, "Analyzing Error-Prone System Structure," *IEEE Trans Software Eng* 17, 2, Febr 1991, pp. 141–152.

[23] J. M. Wing, "A Specifier's Introduction to Formal Methods," *Computer*, Vol. 23, Sept 1990, pp. 8–26.