

AR
93
125-III, 1

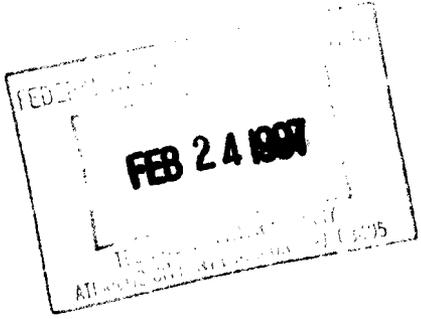
COPY - 2

DOT/FAA/AR-95/125-III, 1

Office of Aviation Research
Washington, D.C. 20591

Digital Systems Validation Handbook, Volume III

Formal Methods and Their Role in Digital Systems Validation for Airborne Systems, Chapter 1



January 1997

Final Report

This document is available to the U.S. public through the National Technical Information Service, Springfield, Virginia 22161



U.S. Department of Transportation
Federal Aviation Administration



NOTICE

This document is disseminated under the sponsorship of the U.S. Department of Transportation in the interest of information exchange. The United States Government assumes no liability for the contents or use thereof. The United States Government does not endorse products or manufacturers. Trade or manufacturer's names appear herein solely because they are considered essential to the objective of this report.

Rushby, J.
Digital systems validation
handbook Volume III...



1. Report No. DOT/FAA/AR-95/125-III, 1	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle DIGITAL SYSTEMS VALIDATION HANDBOOK VOLUME III, FORMAL METHODS AND THEIR ROLE IN DIGITAL SYSTEMS VALIDATION FOR AIRBORNE SYSTEMS, CHAPTER 1		5. Report Date January 1997	
7. Author(s) J. Rushby, SRI International		6. Performing Organization Code	
9. Performing Organization Name and Address SRI International Computer Science Laboratory Menlo Park, CA 94025		8. Performing Organization Report No.	
12. Sponsoring Agency Name and Address U.S. Department of Transportation Federal Aviation Administration Office of Aviation Research Washington, DC 20591		10. Work Unit No. (TRAIS)	
15. Supplementary Notes Peter J. Saraceni, William J. Hughes Technical Center Program Manager, (609) 485-5577, Fax x-4005, saracenp@admin.tc.faa.gov		11. Contract or Grant No. DTFA03-89-C-00043	
16. Abstract The purpose of this Handbook Chapter is to explain the use of formal methods in the specification and verification of software and hardware requirements, designs, and implementations, to identify the benefits, weaknesses, and difficulties in applying these methods to digital systems used in airborne systems, and to suggest factors for consideration when formal methods are offered in support of certification. The presentation concentrates on the rationale for formal methods and on their contribution to assurance for airborne systems within the context provided by Requirements and Technical Concepts for Aviation's DO-178B. This chapter is intended as an introduction for those to whom these topics are new. A more technical discussion of formal methods is provided in NASA Contractor Report 4551, "Formal Methods and Digital Systems Validation for Airborne Systems."		13. Type of Report and Period Covered Tutorial Handbook Chapter 1	
17. Key Words Formal methods, Formal specification, Correctness, Theorem proving, Proof checking, Abstraction, Design faults, Formal logic, Formal proof, Formal deduction, Formal validation, Formal verification, Certification		18. Distribution Statement This document is available to the U. S. public through the National Technical Information Service, Springfield, Virginia 22161.	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 52	22. Price

TABLE OF CONTENTS

<u>Section</u>	<u>Page</u>
EXECUTIVE SUMMARY	1-vii
1. INTRODUCTION	1-1
2. THE RATIONALE FOR FORMAL METHODS	1-1
2.1 The Problem with Software and Its Assurance	1-1
2.1.1 Complexity and Design Faults	1-1
2.1.2 The Discontinuous Behavior of Software Systems	1-5
2.2 Formal Methods	1-7
2.2.1 Analytic Formal Methods	1-7
2.2.2 Descriptive Formal Methods	1-10
3. ISSUES AND CHOICES IN FORMAL METHODS	1-11
3.1 Selection and Abstraction in Applications of Formal Methods	1-12
3.1.1 Levels of Formality	1-13
3.1.2 Selected Components	1-17
3.1.3 Selected Properties	1-18
3.1.4 Life Cycle Stages	1-18
3.1.5 Abstraction	1-19
3.2 The Varieties of Formal Specifications	1-20
3.2.1 Model-Oriented Specifications	1-21
3.2.2 Property-Oriented Specifications	1-22
3.2.3 Specifications for Concurrent Systems	1-23
3.3 The Varieties of Formal Analysis	1-23
3.3.1 Consistency Analysis and Typechecking	1-23
3.3.2 Validating Formal Specifications	1-24
3.3.3 Predicting Behavior and Verifying Refinement	1-26

4. FORMAL METHODS AND CERTIFICATION	1-28
4.1 General Recommendations	1-28
4.2 Interpretation for RTCA/DO-178B	1-29
5. SOFTWARE VERIFICATION PROCESS	1-32
6. CONCLUSION	1-37
7. BIBLIOGRAPHY	1-38
8. GLOSSARY	1-42

LIST OF ACRONYMS AND ABBREVIATIONS

FTA	Fault-Tree Analysis
FMECA	Failure Modes, Effects, and Criticality Analysis
DOT	Department of Transportation
FAA	Federal Aviation Administration
ASIC	Application-Specific Integrated Circuit
AACS	Attitude and Articulation Control System
CDS	Command and Data Subsystem
RTCA	Requirements and Technical Concepts for Aviation, Inc.

EXECUTIVE SUMMARY

The purpose of this handbook chapter is to explain the use of formal methods in the specification and verification of software and hardware requirements, designs, and implementations; to identify the benefits, weaknesses, and difficulties in applying these methods to digital systems used in airborne systems; and to suggest factors for consideration when formal methods are offered in support of certification. The presentation concentrates on the rationale for formal methods and on their contribution to assurance for airborne systems within the context provided by Requirements and Technical Concepts for Aviation's DO-178B. This chapter is intended as an introduction for those to whom these topics are new. A more technical discussion of formal methods is provided in NASA Contractor Report 4551, "Formal Methods and Digital Systems Validation for Airborne Systems."

1. INTRODUCTION.

The purpose of this chapter is to outline what is meant by “formal methods” and to explain their rationale and suggest techniques for their use in providing assurance for airborne systems in the context of RTCA/DO-178B (1992). This chapter is intended as an introduction for those to whom these topics are new and assumes no background beyond some exposure to software engineering and to safety-critical systems. A more technical examination of formal methods is provided in a companion report (Rushby 1993).

The formal methods portion of this presentation is in three sections: section 2 outlines the general rationale for formal methods; section 3 considers the different kinds of formal methods, and some of the issues in their selection and application; and section 4 considers their contribution to assurance for airborne systems, paying special attention to the requirements of RTCA/DO-178B.

2. THE RATIONALE FOR FORMAL METHODS.

Formal methods are a very different approach to software development and assurance than traditional methods. In order to describe why formal methods can be worthwhile, it is necessary to explain why the assurance problem is so hard for software.

2.1 THE PROBLEM WITH SOFTWARE AND ITS ASSURANCE.

Software is notorious for being late, expensive, and wrong. Exasperated technical managers often ask “what is so different about software engineering; why can’t we (or, less generously, you) do it right?” The unstated implication is that the traditional engineering disciplines in which technical managers usually received their training do things better.

This unflattering comparison of software with other engineering endeavors is somewhat justified; in particular, the traditional disciplines are founded on science and mathematics and are able to model and predict the characteristics and properties of their designs quite accurately; whereas software engineering is more of a craft activity based on trial and error rather than calculation and prediction. However, the comparison is too glib in that it fails to acknowledge that in two important respects software is different. These respects are the complexity of behavior that is achieved by software and its lack of continuity. These are discussed in the next two sections.

2.1.1 Complexity and Design Faults.

Software provides much of the functionality of modern systems, and software therefore directly expresses the scale and complexity of these systems. Complexity is a source of design faults, which are faults in the intellectual construction of the system faults that will cause the system to do the wrong thing in some circumstances. Design faults can occur in any system independently of the technologies used in its construction (see, for example, Brazendale and Jeffs 1994) but, because design faults are often due to a failure to anticipate certain interactions among the components of the system or between the system and its environment, they become more likely as the number and complexity of possible behaviors and interactions increase.

Individual software components perform complex functions in modern systems, and collectively they provide the focus for interaction among all parts of the system and between the system and its environment and operators (i.e., pilots). Furthermore, software, because of its mutability, is also the target for most of the changes that are generated in requirements and constraints as the overall design for a system evolves. Thus, software carries the burden of overall system complexity and volatility, and it is to be expected that design faults will most commonly be expressed in software.

2.1.1.1 The Need for Correctness.

Because software is found in active control systems, it is usually infeasible to compensate for possible faults or uncertainties in its design by “overengineering” it to provide a “design margin” in the same way as physical systems. Whereas, a wing spar may be constructed to withstand loads far greater than any it should encounter in normal flight, the software in an autoland system, for example, has to do exactly the right thing.

NOTE: A plausible counterpart to over-engineering in software may be defensive programming, whereby each software component explicitly checks for “impossible” conditions and tries to do something sensible if they arise. The problem is that if the impossible happens, then some failure of design must have already occurred, and there is no telling what impact an autonomous decision to do something locally “sensible” may have on overall system behavior. This is the central problem with complex, interacting systems: local actions can have highly nonlocal consequences.

Another technique that does not extend from physical to design faults is simple modular redundancy. There is always the possibility that physical components may fail, so it is a good idea to have spares and backups to provide fault tolerance. A fault-tolerant system must be designed to avoid common mode failures in which all its redundant components are brought down by a single cause. An example of a common-mode failure is the loss of all hydraulic systems in the Sioux City DC-10 crash: the separate hydraulic systems were routed through a single space near the tail engine and all were severed when that engine disintegrated.

Design faults are the quintessential source of common mode failures, so simple replication can provide no protection against them. It is possible to provide redundant components based on different designs, so-called multiple-version dissimilar (or diverse) software but this is not a fully satisfactory solution. A very brief summary why this is so is given in a note on page 1-3. The general topic of fault-tolerant software (that is, software that can tolerate faults in its own design) is covered in detail in (Hecht 1989).

2.1.1.2 Evidence for Correctness and the Need to Consider All Behaviors.

Although defensive programming and software diversity provide palliatives in some circumstances, for most critical software systems, there is no alternative to the daunting task of eliminating all design faults or at least those that could have serious consequences. It is also necessary to provide evidence that this has been done successfully. This evidence is usually in two parts: one, which is concerned with the process of design and construction, seeks to show through evidence of good practice that everything has been done to prevent serious design faults being introduced and

remaining undetected and uneradicated; and the second seeks to demonstrate directly, through examination of the system in operation and under test and through an analysis of its design and supporting rationale, that it is free of serious faults. The first of these forms of evidence concerns quality control; the second provides quality assurance.

Assurance for a safety-critical system must, at least in principle, consider all possible behaviors of the system under all the circumstances it might encounter. Since “all possible” behaviors may be too many to examine, two complementary approaches have evolved that attempt to reduce the number of behaviors that must be considered. One way tries to show that the system always does the right thing; the other tries to show that it never does a seriously wrong thing.

NOTE: The topic of constructing systems that can tolerate faults in their own design using multiple-version software is controversial. The main questions are whether this approach provides any significant additional assurance of safety and whether that assurance is quantifiable. Answers to these questions hinge on “how much” dissimilarity of design can be achieved in the different versions and on the extent to which any failures of dissimilar designs will be independent.

The extent of dissimilarity depends on how much of the overall design is developed in multiple versions. If dissimilar design is limited to multiple versions of low-level modules, then no protection is provided against design faults above that level; in particular, the system is fully exposed to faults in the modules’ requirements. Furthermore, the degree of design freedom, and hence the scope for diversity, is limited when small components are built to a common set of requirements, and there is some evidence that different designers or implementers do tend to make similar mistakes (Eckhardt et al. 1991, Knight and Leveson 1986).

If dissimilarity is at the level of whole systems or subsystems (e.g., an independent backup to a digital flight control system), then there is the question whether the dissimilar system should have the full capability and assurance of the primary system: if it does, then development and maintenance costs will be at least doubled (and that money could have been spent improving the quality, or the assurance, of the primary system); if not, there is concern whether the secondary system can be relied on in an emergency (e.g., the control envelope of an analog backup system is often less than that of the primary flight control system).

In all cases, there is the critical problem of designing and implementing redundancy management across the dissimilar versions: that is, how to decide when one version has failed and another should be given control (in the case of backup systems), or how to resolve voter disagreements in the case of parallel systems. (Dissimilar designs cannot be expected to produce bit-for-bit identical behavior, so threshold voting has to be used.) Like other problems involving synchronization and coordination of concurrently active distributed components, redundancy management whether of identical or dissimilar components is among the most difficult and fault-prone aspects of software design. Redundancy management does

not lend itself to diversity (e.g., one cannot vote the voters ad infinitum; ultimately a decision must be made and the algorithm by which that decision is accomplished represents a single design) and can be made more complex and fault prone by the need to manage diversity in other components. For example, on test flight 44, disagreements among the threshold voters in the AFTI-F16 digital flight control system caused each computer to declare the others failed. The analog backup was not selected because simultaneous failure of two or more digital channels had not been anticipated in design of the redundancy management system (Mackall 1988).

For these and other reasons, the guidelines for certification of airborne software state that the degree of protection provided by software diversity “is not usually measurable” and dissimilar software versions do not provide a means for achieving safety-critical requirements; however, “are usually used as a means of providing additional protection after the software verification process objectives for the software level...have been met” (RTCA/DO-178B 1992).

For the first approach, a combination of analysis and empirical testing to examine those behaviors that are considered most likely to harbor serious faults are used. For example, those that are close to boundary conditions or that represent “off nominal” conditions, such as those where some subsystems or redundant components have failed. Fault injection (an empirical method) and failure modes, effects and criticality analysis (FMECA), an analytical method, are some other examples.

The general idea behind the second approach to quality assurance is to hypothesize that the system has done something bad and then to analyze all the circumstances that could cause this to come about and to show that the design prevents them from happening. This approach is inspired by hazard analysis, which is a central concept in safety-critical systems; one particular method for doing it that has been adapted to software is fault-tree analysis (FTA).

The property that is common to the different assurance techniques is that they provide ways to group “essentially similar” behaviors together so that fewer cases need to be considered while still providing effectively complete coverage of all possible behaviors. These techniques are very effective with systems based on mechanical, hydraulic, electrical, and other physical components: these have relatively few “essentially different” behaviors, so that relatively straightforward analysis combined with a modest number of empirical tests is sufficient to cover all cases. These familiar techniques are far less effective, however, with complex systems that can exhibit extremely large numbers of essentially different behaviors.

Because the complexity in modern systems is expressed in software, it follows that the software will exhibit a large number of different behaviors and that assurance will be difficult for this reason. In fact, this difficulty is compounded by another attribute of software that distinguishes it from physical systems. This attribute is considered next.

2.1.2 The Discontinuous Behavior of Software Systems.

The reason that software is the focus for most of the design complexity in modern systems is its versatility: a software system can provide many different behaviors and can be programmed to respond appropriately to many different circumstances. The source of these different behaviors and responses is in the many discrete decisions that are made as software executes: each decision is discrete in that the subsequent course of execution switches from one path to another according to whether or not some condition is true. Because the relationship between the inputs and outputs of a piece of software is the cumulative effect of these many discrete decisions, it follows that the overall input/output relationship must itself be discretized or discontinuous: small changes in inputs can change the outcomes at certain decision points, resulting in radically changed execution paths and correspondingly large changes in output behavior. This discontinuous relationship between inputs and outputs is the second major respect in which software differs from the physical processes considered by other engineering disciplines.

In physical systems, there is usually a (piecewise) continuous relationship between inputs and outputs: smooth changes in the inputs produce correspondingly smooth changes in the outputs. This allows the complete behavior of a physical system to be extrapolated from a finite number of tests: the continuous character of the system ensures that responses to untested input configurations will be essentially similar to those of nearby cases that have been tested. Departures from continuity are usually catastrophic breakdowns in response to inputs beyond the operating range.

However, with software, this method of inferring properties of the totality of possible behaviors from tests on a selected sample is much less secure: without continuity, it cannot be assumed that neighboring cases are essentially similar to one another, so there is little justification for extrapolating from tested to untested cases. Although less than exhaustive testing does not allow definitive statements to be made about complex software, it does permit statistical statements of its reliability and that such quantification of reliability is both necessary and sufficient for the certification of safety-critical systems. Sometimes this is countered by the argument that talk of reliability is meaningless when dealing with design faults. If design faults are present, they will cause the system to fail in specific circumstances, and the failure is certain whenever those circumstances arise. However, it must be recognized that occurrence of those circumstances is associated with a random process, namely, the sequence of inputs to the system (or, more generally, the sequence and timing of its interactions with its environment). Thus, the manifestations of design faults behave as stochastic processes and can be treated probabilistically. A piece of software having a failure rate of less than 10^{-9} per hour is to say that the probability of encountering a sequence of inputs that will cause a design fault to lead to failure is less than 10^{-9} per hour.

NOTE: Although this chapter speaks only of software, exactly the same concerns apply to many hardware components, especially custom application-specific integrated circuits (ASICs). These share all the important properties of software, notably, design complexity and discontinuity of behavior and differ only in the technology of their implementation. Whereas software design is ultimately expressed in a programming language such as Ada and then compiled into code that is interpreted by a processor, ASIC designs are expressed in a hardware design

language such as VHDL or Verilog and then transformed into hardware structures or to gate-array configurations. The considerations for assurance described in standards and guidelines such as RTCA/DO-178B apply to ASICs as they do to software. Similarly, the techniques of formal methods can be applied to ASICs and other complex hardware designs.

The problem with the experimental statistical approach to assurance for complex software is that the smallest failure rates that can be determined in this way are typically several orders of magnitude greater than those required for safety-critical systems. This is explained in somewhat more detail in the next NOTE.

The infeasibility of experimental quantification of reliability for safety-critical software means that its assurance must chiefly be provided by other means. Experimental evaluation is not the only means for providing assurance about the behaviors of physically engineered systems. The engineering field concerned normally provides well-validated mathematical models that allow the properties and behavior of a given design to be predicted through calculation. For example, structural engineers can calculate the behavior of a wing spar before it is built. It is one of the distinctions between engineering and craft activities that engineering uses mathematical modeling to predict behavior; whereas crafts use trial and error. Formal methods, which are introduced in the next section, provide a way to move the construction and validation of software away from experiment and adjustment and toward prediction and calculation.

NOTE: It is perfectly reasonable to state requirements for safety-critical systems in statistical terms. For example, catastrophic failure conditions in aircraft (“those which would prevent continued safe flight and landing”) must be “extremely improbable.” That is, “so unlikely that they are not anticipated to occur during the entire operational life of all airplanes of one type” (Advisory Circular 25.1309-1A, paragraphs 6.h[3] and 9.e[3]). A little arithmetic suggests 10^7 hours as the operational lifetime of an aircraft fleet, and hazard analysis might typically reveal ten potentially catastrophic failure conditions in each of ten systems on-board the aircraft, so that the maximum allowable failure rate for each is about 10^{-9} per hour (Lloyd and Tye 1982). This is indeed the number suggested as an “aid to engineering judgment to help determine compliance” with the requirement for extremely improbable failure conditions (Advisory Circular 25.1309-1A, paragraph 10.b). The probability 10^{-9} is applied to complete (sub)system failure, not to any software the system may contain. Numerical estimates of reliability are not assigned to airborne software (RTCA/DO-178B 1992), but the figure gives an idea of the quality required.

For a simple physical system where breakdown or wearout is the only potential cause for a catastrophic failure condition, experience with similar systems together with testing and analysis may yield data that can substantiate a claimed failure rate as low as 10^{-9} .

Software generally interacts with its environment in such a complex manner that prior experience of its behavior in similar applications may provide relatively little assurance in a new one. In flight tests of the X31, the control system “went into a reversionary mode four times in the first nine flights, usually due to disagreement between the two air-data sources. The air-data logic dates back to the mid-1960s and had a divide by zero that occurred briefly. This was not a problem in its previous application, but the X31 flight-control system would not tolerate it” (Dornheim 1991). Similarly, much of the software in the Therac 25 medical electron accelerator, which led to massive overdoses of radiation and the subsequent deaths of six patients, had been used in an earlier machine without accident (Leveson and Turner 1993). Furthermore, most software has significant elements of novelty from one application or version to another, so that experimental determination of software reliability must examine the actual software in the context of its actual application. Furthermore, the test scenarios used to derive reliability estimates must closely approximate in type and frequency the distribution of inputs that will be encountered in operation (this is called the operational profile). For required failure rates on the order of 10^{-9} , this means that it will be necessary to construct many millions of the very rare scenarios that will each be encountered only one time in a billion. Catastrophic failures usually arise in situations compounded by several rare events (Hecht 1993). Divergence between the test and operational profiles in these remote regions can lead to inaccurate estimates of reliability and spurious assurance of safety.

The difficulty in reproducing the operational profile for rare events and the time required to perform fault injections and to configure other elements of “all-up” test scenarios limit the feasible failure rates that can be determined empirically to 10^{-4} or 10^{-5} (Butler and Finelli 1993) nowhere near the 10^{-9} that is required.

2.2 FORMAL METHODS.

The term formal methods refers to the use of mathematical modeling, calculation, and prediction in the specification, design, analysis, construction, and assurance of computer systems and software. The reason it is called “formal methods” rather than “mathematical modeling of software” is to highlight the character of the mathematics involved.

2.2.1 Analytic Formal Methods.

Each engineering discipline develops a body of mathematical techniques that are particularly appropriate for modeling and predicting the phenomena relevant to its field. In many cases, the relevant applied mathematics uses partial differential equations to model the variations in continuous physical quantities over time or space. For software, the familiar methods of calculus and differential equations are inapplicable because, as noted above, discrete, rather than continuous quantities, have to be modeled. Care must be taken here to distinguish the mathematics of the domain to which the software is applied (which may, as in the case of control applications, require the evaluation of expressions derived from differential equations) from the

mathematics that describes its own operation. Instead of differential equations, the properties and behaviors of concern are best described in terms of concepts from discrete mathematics: sets, graphs, partial orders, finite-state machines, and so on. Calculation in these finite domains is based on the methods of formal (or mathematical) logic rather than numerical computation. This is because the results of interest are logical properties, such as “this system can tolerate any single fault in any component,” rather than numerical estimates for some parameter such as lift or drag.

To deduce whether a certain logical property follows from descriptions of certain discrete mathematical structures, one has to start from the axioms describing those structures and manipulate their symbols according to certain rules of deduction. This process is more akin to proving theorems in Euclidean geometry than to ordinary numerical calculation, but it shares with calculation the characteristic that it is performed according to strict rules, so that one person can check the work of another and computers can be used to automate some of the steps. The process of manipulating symbols according to certain rules is called “formal deduction” because the legitimacy of the process depends only on the form of the symbolic expressions concerned and not on what they are supposed to mean. For example, the transformation of an expression of the form $x^2 - y^2$ into one of the form $(x + y) \times (x - y)$ is legitimate whenever x and y are numbers, independently of whether they represent the mass of planets or the debts of nations.

The particular importance of formal methods to safety-critical systems is that they, though subject to caveats discussed shortly, permit analysis of all the behaviors of a software system. This total exploration is the only way to provide assurance that catastrophic failure does not lie hidden among the vast number of possible behaviors. How is total exploration possible? It has already been seen that in the absence of continuity there is no basis for extrapolating from tested to untested cases, so how is it that a finite procedure based on formal methods can provide assurance for all the (possibly infinite) behaviors of a software system?

Part of the explanation is that formal methods provide powerful tools for identifying and grouping essentially similar pieces of behavior together so that all members of a group can be dealt with at a single shot. In empirical testing, only the external manifestations of the system are examined, and the ability to assign these to groups that are essentially similar is very limited because of their discontinuous nature. However, in formal methods the internal design of the system is examined, where the sources of discontinuity are visible, and little pieces of behavior can be grouped together (all those that result from following a certain path through a certain part of the design, for example). The properties of those pieces can be characterized by mathematical expressions (i.e., formal specifications) and can deduce the properties of larger pieces of behavior by applying formal deduction to the expressions describing their component pieces. By composing small pieces of behavior together to yield larger and larger parts of the complete behavior, all possible end-to-end behaviors are eventually covered without having to enumerate them explicitly.

Another way of saying this is that formal methods allow the calculation of the reasons why a software design does its job. Software is a designed artifact consciously constructed to achieve a goal: if the goal can be written down, and how the design accomplishes it, then an argument could be constructed that explains why the software does its job. Formal methods allow this argument to be reduced to the certainty of calculation.

The certainty of calculation needs qualification: for example, the rules of arithmetic reduce addition to the certainty of calculation, but it is easy to make mistakes when adding a long column of figures. The calculations underlying formal methods are similarly tedious and error prone when done by hand, so it is often desirable to automate them. Unfortunately, these kinds of calculations are not so easily mechanized as numerical ones. Whereas the steps of a numerical calculation can be programmed as a deterministic algorithm, selection of the steps in a formal deduction requires either insight, or a heuristically guided trial-and-error search. Theorem provers are computer programs that attempt to automate formal deductions through a combination of heuristics and brute-force search; proof checkers are programs that leave selection of the steps to an insightful human and simply check that each one is carried out correctly. The most effective automated reasoning tools for formal methods generally combine elements from both theorem provers and proof checkers. For simplicity, the term “theorem prover” is generally used in the following sections to cover all forms of automated deduction.

NOTE: Mathematical logic provides ways to reason about the properties of large or infinite collections of related things in a finite manner: instead of reasoning about the behavior of a system when given the input 1, or the input 2, or..., logic provides methods to reason about its behavior on the symbolic input (n) thereby collapsing all the separate cases into one. One method allows conclusions to be drawn for all values of a given variable by reasoning about a single representative symbolic constant (this process is called Solemnization). Another method allows the deducing of properties for all values of some ordered domain (e.g., the natural numbers, 0, 1, 2...) by showing (a) that the property is true for the least element(s) of the domain (e.g., 0) and (b) that when it is true for all members up to some point (e.g., n) then it is also true for the next point (e.g., n + 1) (this process is called mathematical induction). Formal deduction allows properties of a complete system to be deduced by combining basic steps such as these: steps that are required to follow certain rules and can therefore be checked by others or by machine.

Earlier, it was stated that the ability to represent values symbolically, and hence to group related behaviors together, is part of the reason why formal methods allow the consideration of all the vast numbers of behaviors of a complex software system, and it was also stated that this claim is subject to certain caveats. Both points concern the fact that formal methods are a modeling activity: formal methods do not deal with actual software running on electronic computers interacting with the real environment but with mathematical models of these artifacts. This is exactly similar to the mathematical methods used in other engineering disciplines: a finite-element calculation does not calculate the stress in wing spar; it calculates a representation of the stress in a mathematical model of a wing spar.

Modeling is a source of both weakness and strength. Because a model is not the same as reality, predictions made with its aid may be incorrect. This can be because the model is insufficiently detailed, or because it is wrong. It is necessary to guard against these potential weaknesses by carefully validating models before trusting their predictions. The fear of imprecision or inaccuracy sometimes leads to the development of complicated, highly detailed models, but this can vitiate the main purpose for developing a model in the first place: its ability to support tractable analysis. The

great opportunity offered by modeling is the freedom to select and simplify the aspects of reality that are to be considered; Newtonian mechanics achieves its effectiveness because it selects for attention just a few properties of bodies and their motions: the mass of a body is considered, but not its volume or its color. In the same way, formal methods can achieve great effectiveness by focusing on just certain parts of a complex system (e.g., those that give the most difficulty, such as redundancy management) and by excluding details that are judged irrelevant (e.g., the algorithms for redundancy management may be focused on, and the details of their execution as programs could be ignored). It requires great skill and judgment to perform the abstraction necessary to create formal models that are simple enough to be computationally tractable, yet realistic enough to provide accurate predictions and credible assurance.

2.2.2 Descriptive Formal Methods.

In the previous discussion, formal methods were introduced by analogy with engineering mathematics, and it was stated that the purpose of formal methods is to make predictions about the properties and behavior of software based on calculations performed on a mathematical model of its design. However, engineering mathematics is not used only for calculation and analysis, it also provides a vocabulary for describing and documenting designs and a framework for thinking about them. Thus, aeronautical engineers may speak of drag divergence or flow separation as concepts independently of particular sets of equations. Formal methods can serve a similar descriptive function for software. Concepts such as relations, functions, finite-state machines, and universal quantification can supplement or replace the English prose, pseudocode, and various kinds of diagrams that are traditionally used in documenting requirements, specifications, and designs for software systems. Two benefits can follow from this use of mathematical notation: it can improve the quality of documentation and lead to better communication among those working on the system, and it can supply better ways of thinking about software.

There are two aspects to a software system: control and data. Control is concerned with the selection, timing, and sequencing of the operations performed in a software system. Data is concerned with how information is represented within the computer system and manipulated by its software. The problem with traditional ways of documenting and thinking about both aspects of software is that they are almost entirely operational: software is understood by mentally executing it; similarly specifications all too often describe how the software works rather than what goals it is to accomplish.

For the control aspect, traditional operational methods can be quite effective for systems composed of a single sequential program: with practice and determination, it is often possible to think through the consequences of each alternative at a branch point, and the behavior of loops can largely be understood by mentally considering the cases where each is executed zero, one, and many times. For large or reactive or parallel systems, however, these methods become very unreliable. It is hard to comprehend all the possibilities when external events can have an impact at almost any point and scenarios such as “what happens if a timer interrupt occurs here, and suppose the other processor is in fault-recovery mode, so that it could post a ‘need-service’ flag at just the instant that...” have to be considered. The note on page 1-11 describes an example of the difficulties caused when one of the possible scenarios is overlooked.

As with control, the traditional way of documenting and thinking about the data aspect of software is largely in terms of its concrete representations. Thus, much software documentation is composed of pictures or programming-like notation describing how certain information is recorded in the bits of a computer word, how collections of similar items are represented in arrays of computer words, and how more complex structures are represented by pointers linking records together. More modern notations derived from programming languages like Ada and ideas from object-oriented design (Rumbaugh et al. 1991), have significantly raised the level of abstraction and improved the organization of data descriptions, but the orientation is still that of implementation.

In contrast to these operational ways of describing and understanding software systems, formal methods provide ways to document and think about data and control that depend less on mentally tracing execution paths and more on identifying properties that are to be assumed, established, or preserved. In addition, the mathematical concepts employed in formal methods assist in the construction of more abstract descriptions that state more clearly what is to be accomplished without becoming caught up in premature details of how it is to be done.

Although the main purpose of descriptive formal methods is to improve communications among those working on a software project and to facilitate informal quality control activities such as reviews and inspections, these methods are not antithetic to the analytical purposes described earlier. Rather, there is a continuum of formal methods ranging from those that are primarily intended to support mechanically checked analysis to those that are primarily intended for documentation.

3. ISSUES AND CHOICES IN FORMAL METHODS.

A sketch has been given of the rationale for using formal methods in software development and assurance. In this section, some of the different varieties of formal methods are described together with some considerations for their selection and use.

NOTE: While in orbit around Venus, the Magellan spacecraft broke contact with Earth and entered "safing" modes preempting its scientific mission on a number of occasions. Extensive efforts were undertaken to find the source of the problem, and after eight months, the most likely cause was identified (Kasuda and Packard 1993). Two flags determine whether a background task should be run in the otherwise unused time after the end of all the foreground tasks in the current frame and before the end-of-frame interrupt. The scheduled flag determines whether a particular task should be run as the background task, and the active flag indicates whether this task is an uncompleted activity that should be resumed at an address stored on the stack or a new one that should start at its entry point. On very rare occasions, the end-of-frame timer interrupt would occur in the instant after one flag had been set to a new value but before the other had been. In particular, a sequence that was invoked when a background task completed could be interrupted after the scheduled flag had been turned off but before the active flag could be turned off also. Next time this task was scheduled, the background task manager would mistakenly think it was

active and would pop the stack to obtain the address at which to continue its execution. Since the task had completed, there was no restart address on the stack, so the value that was popped and used was some random piece of data. This random address sent the processor to a piece of code where it sat in a tight loop that continually reset the watchdog timer, thereby disabling the very mechanism that was intended to thwart such runaways (Cooper 1993). Fortunately, Magellan had multiple levels of redundancy, and although these were intended to cope with hardware, not design problems, they saved the spacecraft. Specifically, the design fault described above was in the software of the attitude and articulation control system (AACS) computer and although the runaway execution reset the watchdog timer it did not modulate the heartbeat pattern that is placed in memory shared with the command and data subsystem (CDS) computer. When the CDS computer saw the AACS heartbeat cease, it reconfigured the AACS systems, eventually leading to a reboot of the errant AACS computer.

Computer scientists are thoroughly familiar with the dangers of being interrupted while adjusting critical data structures and will normally arrange for such actions to take place inside a critical section that cannot be interrupted. These incur overhead, however, and the Magellan designers thought that in their particular circumstance it was safe to do without this protection. Cooper's book, cited above, aptly conveys the monumental task of trying to diagnose a very rare but devastating misbehavior in a reactive, real-time, parallel software system operating in the presence of faults when the only way to understand the system is to mentally (or actually) simulate its execution under as many circumstances as possible.

3.1 SELECTION AND ABSTRACTION IN APPLICATIONS OF FORMAL METHODS.

Expertise in formal methods is not widespread and can be costly to acquire. Furthermore, the resources available for any project are limited, so that effort expended on formal methods may reduce that available for other methods of analysis and assurance. For these reasons, formal methods need to be applied selectively. There are several dimensions in the use of formal methods that permit selective or partial application. Five of the most important are listed below:

- a. The amount of formality can vary between occasional use of ideas and notation from discrete mathematics in a pencil and paper manner to fully formal treatments that are checked with a mechanical theorem prover.
- b. Formal methods can be applied to all or only to selected components of the system.
- c. Formal methods can be applied to selected properties of the system (e.g., absence of deadlock) rather than to its full functionality.
- d. Formal methods can be applied to all or merely to some of the stages of the development life cycle. If the latter, a choice should be made of whether to favor the earlier or the later stages of the life cycle.

e. In all cases it is possible to include more or less detail and to choose the level of abstraction at which the formal treatment is conducted.

Each of these is examined in more detail below.

3.1.1 Levels of Formality.

The very notion *formal* can have different interpretations, and methods differ in the amount of formality they employ. For example, suppose a program is produced that will compute the exponential n^m where n and m are integers and m is nonnegative. One way to do this is by repeated multiplication; it is not a very efficient method but will serve the purpose at hand. The following program, written in a generic high-level language, computes the value n^m and leaves it in the variable r :

```
r: = 1;
i: = m;
while i ≠ 0 do
  r: = r * n;
  i: = i - 1
endwhile
```

The following proof justifies the claim that the program satisfies its specification.

a. Each time in the while loop, at the point just before the $i \neq 0$ test, the following relationship is true among the variables:

$$r \times n^i = n^m.$$

To prove this, consider two cases:

1. First time into the loop. The following have been initialized: r to 1 and i to m , so $r \times n^i$ is $1 \times n^m$ and the desired relationship is true.

2. Other times around the loop. Assume the previous iteration left $r \times n^i = n^m$; after going round the loop once more, r has been replaced by $r \times n$ and i by $i - 1$, so it needs to be proved that $(r \times n) \times n^{i-1} = n^m$. By arithmetic, the left side equals $r \times n^i$ and the result follows.

b. The while loop is exited when $i = 0$. Since it is known that $r \times n^i = n^m$ at this point, it follows by arithmetic that n^i is 1, and so $r = n^m$ as required.

c. To see that the program always terminates, note that i is initialized to m , which is a nonnegative integer. The value of i is reduced by one each time around the loop, so eventually it will reach zero and the loop will exit.

Although a proof was used here, the process was not particularly formal, the argument was presented in fairly ordinary English, and relied on intuitive understanding of how the program executes.

A more formal approach would use logical axioms to describe the behavior of the program without requiring one to mentally execute it. The following example uses a method that manipulates Hoare sentences, which are constructs of the form $\{P\}S\{Q\}$, where P and Q are expressions describing the relationships among the program variables, and S is a piece of program text. The interpretation is that if the relationship P is true before execution of S , and if S terminates, then the relationship Q will be true afterward.

The behavior of a while loop is specified by the following axiom. (In this formula, the symbol \wedge means “and,” and \neg means “not”.)

$$\frac{\{P \wedge B\}S\{P\}}{\{P\} \text{while } B \text{ do } S \text{ endwhile } \{P \wedge \neg B\}}$$

This shows if the Hoare sentence above the line is true then the Hoare sentence below the line will be true also. If $i \neq 0$ is substituted for B , $r := r * n; i := i - 1$ for S , and $r \times n^i = n^m$ for P , then the following is obtained:

$$\frac{\{r \times n^i = n^m \wedge i \neq 0\} r := r * n; i := i - 1 \{r \times n^i = n^m\}}{\{r \times n^i = n^m\} \text{ while } i \neq 0 \text{ do } r := r * n; i := i - 1 \text{ endwhile } \{r \times n^i = n^m \wedge i = 0\}}$$

Now the expression $r \times n^i = n^m \wedge i = 0$ at the bottom right gives $r = n^m$, so the next step is to prove the Hoare sentence above the line, which involves the sequential composition of two assignment statements.

The axiom for sequential composition is

$$\frac{\{P\}S_1\{Q\} \wedge \{Q\}S_2\{R\}}{\{P\}S_1; S_2\{R\}}$$

and $r \times n^i = n^m \wedge i \neq 0$ for P , $r := r * n$ for S_1 , $r \times n^i = n \times n^m$ needs to be substituted for Q , $i := i - 1$ for S_2 , and $r \times n^i = n^m$ for R . This gives

$$\frac{\{r \times n^i = n^m \wedge i \neq 0\} r := r * n \{r \times n^i = n \times n^m\} \wedge \{r \times n^i = n \times n^m\} i := i - 1 \{r \times n^i = n^m\}}{\{r \times n^i = n^m \wedge i \neq 0\} r := r * n; i := i - 1 \{r \times n^i = n^m\}}$$

Glossing over some details, the axioms that specify assignment statements allow one to prove the two Hoare sentences above the line. To complete the proof of this program, it needs to be established that the initialization statements establish the relationship assumed at the start of the while loop, that is

$$\{m \geq 0\} r := 1; i := m \{r \times n^i = n^m\}$$

This follows from the axioms for sequential composition and assignment in a manner similar to that of the previous step.

The second way to analyze this program was much more formal than the first; instead of appealing to intuition about how a program executes, axioms and substituted program text and mathematical expressions were used. Although insight and intuition were used to decide which axioms to use and what to substitute into them, the subsequent manipulations were quite mechanical and it should be clear that each of the steps that were performed could have been checked by computer.

There are advantages and disadvantages to the different levels of formality employed in these two treatments of the example program. In discussing them, it is useful to have a simple scale for identifying levels of formality.

Level 1 formal methods use ideas and notation from discrete mathematics and logic but within a loose framework where mathematics, English, diagrams, and other notations are used together. Proofs are careful arguments that are evaluated by whether they persuade reviewers. This is how mathematics is accomplished. The first treatment of the exponentiation program was an example of Level 1 formalism.

Level 2 formal methods employ a fixed specification language for documenting requirements and designs. A specification language generally blends concepts from logic, discrete mathematics, and programming into a single notation. Often, the language is supported by tools that check specifications for certain types of errors and that provide useful functions such as cross-referencing or typesetting. Analyses and proofs are performed by hand and recorded with pencil and paper but make use of explicit axioms and proof rules that describe the semantics of the languages and methods used. The second treatment of the exponentiation program was an example of Level 2 formalism.

Level 3 formal methods stress mechanized analysis. Their specification languages are generally closer to standard logic than those of type 2 formal methods and are supported by tools that include proof checkers, theorem provers, or model checkers. A Level 3 program verification system could mechanize either the first or the second of the approaches presented earlier for the exponentiation program. Most current systems use an approach closer to the first than the second. Typically, a program verification system would require that the program is annotated with assertions and loop invariants as follows:

```

entry assertion  $m \geq 0$ 
r := 1;
i := m;
while i  $\neq$  0 do
  loop invariant  $r \times n^i = n^m$ 
  r := r * n;
  i := i - 1
endwhile
exit assertion  $r = n^m$ 

```

Using the method of inductive assertions, the verification system would then generate the following three verification conditions, which correspond to the paths in the program from the entry assertion to the loop invariant, from the loop invariant around the loop and back to the invariant, and from the loop invariant to the exit assertion, respectively. The symbol \supset in these formulas means “implies.” Also, a prime indicates the value of a variable in the new state following execution of a path, and an unprimed name indicates the value in the old state at the beginning of the path.

- a. VC1: $m \geq 0 \wedge r = 1 \wedge i = m \supset r \times n^i = n^m$.
- b. VC2: $m \geq 0 \wedge r \times n^i = n^m \wedge i \neq 0 \wedge r' = r \times n \wedge i' = i - 1 \supset r' \times n^{i'} = n^m$.
- c. VC3: $m \geq 0 \wedge r \times n^i = n^m \wedge i = 0 \supset r = n^m$.

These verification conditions are expressions in ordinary logic (plus arithmetic) and can be proved quite easily by the theorem provers of most verification systems.

The advantage of Level 1 formal methods is the flexibility that is available: notations and techniques can be selected, or invented, to suit the particular problem at hand. These methods can be very effective when used by individuals or small teams possessing skill and judgment, but the lack of standardized notation and methods can make communication and training difficult across larger groups.

Level 2 formal methods address the problems of communication and training by providing fixed specification notations (Z (Spivey 1993) and VDM (Jones 1990) are well-known examples) and, usually, a methodology for using them. Individual Level 2 methods are well suited to some types of applications (e.g., data processing) and less well suited to others (e.g., concurrent systems); users must be careful not to stretch their chosen method beyond its limits.

In general, the Level 2 notations are optimized for descriptive, rather than analytic, purposes. If the goal is to use formal methods to calculate properties of a design for the purpose of analysis, then a Level 3 method equipped with appropriate tools will probably be more suitable. It generally requires considerable skill and experience to use Level 3 tools effectively, but they can provide a very high degree of assurance.

3.1.2 Selected Components.

Formal methods are generally advocated because it is felt that they can improve quality control and assurance for software. If this is so, then the greatest benefits will be seen when formal methods are applied to the most critical components and to those for which traditional methods have been found least effective.

Subsection 2.2.2 of RTCA/DO-178B identifies software criticality levels A through E according to the severity of their potential failure conditions (i.e., Level A is software whose malfunction could contribute to a catastrophic failure condition). Software criticality level determines the amount of effort and evidence required to show compliance with certification requirements. It provides a natural criterion for selecting components for which formal methods should be considered.

Another criterion that should be considered is the likely effectiveness of formal methods versus traditional methods for quality control and assurance. It is to be expected, and there is some evidence to support the expectation (Lutz 1993), that the intrinsically hard design problems tend to be the most prone to faults and the most resistant to traditional means of assurance. These intrinsically hard problems generally involve complex interactions, such as the coordination of distributed, concurrent, or real-time computations and redundancy management. It requires great skill to address these problems using formal methods, but the number and size of these problems may not be large. Hence, as noted here, the greatest return on formal methods may be obtained when relatively few, very highly skilled people apply formal methods to the hardest and most critical problems.

NOTE: Readers who found their eyes glazing over at the formulas used to verify the trivial exponentiation program may wonder whether these formal techniques really are practical and might ask “How am I going to get my engineers to use this stuff?” Privately, they may also wonder “if this stuff is so good, why isn’t it used more?”

The practicality and cost/benefit of formal methods are heavily dependent on the type of applications considered. Program verification of the kind illustrated in the examples considered here is undeniably tedious and expensive (Guiho and Hennebert 1990) and must compete with traditional methods that are quite effective. It seems that the greatest benefits are likely to be found when formal methods are applied to the hardest and most difficult problems where traditional methods are ineffective or unavailable. Examples of hard problems are those involving distributed and concurrent execution and, especially, redundancy management (Owre et al. 1995). These problems can be considered practical because, though hard, they are not large and can be undertaken by a few highly skilled people. It is not necessary to train every programmer to get valuable returns from formal methods. Another opportunity lies in problems where formal methods can be massively automated. Examples include certain kinds of protocols (Har’El and Kurshan 1990, Clarke et al. 1993) and hardware designs (Miller and Srivas 1995).

Before the slow industrial take-up of formal methods is attributed to ineffectiveness, the comparative youth of the field should be remembered, and the history of other engineering subjects should be recalled (Vincenti 1990). Most of the disciplines now regarded as engineering started as crafts, practiced by experimentalists who learned what would work and what would not by trial and error, and who passed their lore on to their successors by on the job example. Gradually, methods based on science and mathematics started to appear, but they did not immediately displace the traditional methods. Among the reasons for the slow acceptance of mathematical techniques were the conservatism of the traditional practitioners, the arcane difficulty of the new methods (to those lacking the necessary training), and their initially narrow range of application. For example, Tesla quit Edison's laboratory after less than a year complaining of Edison's preference for empirical methods "knowing that a little theory and calculation would have saved him 90% of his labor" (Burns 1993).

The new methods displaced the old as concrete evidence of their superiority accumulated, as good textbooks became available, and as new generations of engineers, trained in the mathematical methods, joined the field. For example, when Donald Douglas, then recently graduated from MIT, started work for the Martin company in 1914, he found Glenn Martin bouncing up and down inside a seaplane supported on wooden trestles "to see if it was strong enough." Douglas noted that "this didn't represent any load that the plane bore in flight." Aided by proper engineering analysis, his first design (the Model S) had almost twice the range and payload of Martin's previous products (Biddle 1991).

It is inconceivable today that an aeronautical engineer could be ignorant of aerodynamics and argue that such mathematical modeling is irrelevant to the practical business of designing airfoils. Formal methods can be expected to play a similar role in software engineering.

3.1.3 Selected Properties.

Just as some components of a system may be more suitable for formal methods than others, so different properties of those components may be more suited than others to formal treatment. As with components, suitability may be determined by criticality or by the effectiveness of formal methods compared with other methods. For example, the important property of a particular component may not be that it does its job (there may be backups to accomplish that), but that it is free of "specific anomalous behaviors" (RTCA/DO-178B 1992). Negative properties such as these (i.e., properties that state what must not happen) are particularly difficult to test and can be good candidates for formal analysis.

3.1.4 Life Cycle Stages.

The example shown earlier involving the exponentiation program illustrated the use of formal methods in the late life cycle. That example was one of program verification, where an executable program is proved to satisfy its detailed requirements. Other applications of formal methods focus

on activities of the early life cycle, such as the documentation and analysis of requirements, and on those of the intermediate life cycle stages such as the documentation of interfaces and the systematic refinement of requirements into designs, or designs into implementations.

Late-life cycle applications of formal methods such as program verification were among the earliest to be developed and are now widely known and well understood. Precisely because this part of the life cycle is well understood, informal methods and engineering practice have achieved a considerable degree of practical effectiveness: sequential programming and gate-level design are not major sources of difficulty or faults today. For example, Lutz (1993) reports on 197 significant software faults detected during integration and system testing of the Voyager and Galileo spacecraft. Only three of these faults were programming errors; the vast majority were requirements problems. Similarly, Keutzer (1991) reports that fully half of all ASICs are faulty on first fabrication and that these faults are invariably due to errors in requirements or high-level design. No errors are reported in implementation below the register-transfer level.

It is now generally recognized that faults introduced in the early life cycle are among the most difficult and expensive to detect and eradicate later. Furthermore, the most serious failures are often traced to undetected faults that were introduced early in the life cycle. One explanation for the intractability and persistence of faults introduced in the early stages of development is that there are few good methods for validating the products of these stages: requirements and early design descriptions do not lend themselves to execution and tests. Formal methods can help overcome this difficulty by allowing early specifications to be challenged and explored through theorem proving: a challenge of the form “if this specification says what it should, then the following ought to follow” can be formulated as a putative theorem that should be provable from the specification. Rapid prototyping can serve some of the same ends, but it is not always straightforward to distinguish those properties that are truly entailed by the requirements or design descriptions being validated from those that are accidental to the prototype. Unlike a prototype, a formal requirements specification can be validated experimentally without necessarily being executable.

3.1.5 Abstraction.

Abstraction is one of the most powerful tools for gaining intellectual mastery of complex systems: it allows one to ignore the irrelevant and simplify the relevant so that the essential matter of concern is exposed to scrutiny in its clearest and most tractable form. Abstraction is also a crucial factor in controlling the size of a formal development and the effort required for its analysis.

One example of abstraction considers the algorithms that underlie a design rather than their expression as programs. For example, the repeated multiplication that underlies the exponentiation program considered earlier can be abstracted to an algorithm represented by the following recursive function. This specification is generally taken as the axiomatic definition of exponentiation and is used to establish lemmas such as $n^1 = n$ and $n^{m1} + m2 = n^{m1} \times n^{m2}$ and the correctness of more efficient algorithms.

$$n^m \stackrel{\text{def}}{=} \text{if } m = 0 \text{ then } 1 \text{ else } n \times n^{m-1} \text{ endif}$$

This specification can be analyzed using ordinary logic (i.e., without the special machinery for program verification, such as Hoare sentences) and it is comparatively easy to establish its properties, even when using Level 3 formal methods and mechanized theorem proving. The transition from this abstract algorithm to the concrete program can be justified either informally, or using formal methods of any of the three levels.

When formal methods are applied to algorithms, there is further scope for abstraction in the choice of how much detail to include. For example, one of the important algorithms in fault-tolerant systems is for distributing sensor samples consistently in the presence of faults (Lamport, Shostak and Pease 1982). This is a distributed algorithm, and if issues of the timing and transfer of the messages that are communicated in the algorithm are of concern, then it is necessary to model these mechanisms in some detail, and the analysis will be correspondingly detailed and lengthy (Lamport and Merz 1994). However, if the main concern is with the fault masking properties of the algorithm, then the mechanisms of distributed computation and communication can be ignored and the algorithm can be modeled as a recursive function, in which form its analysis is quite straightforward. Certain details of behavior, and therefore the opportunity to detect some potential faults, are missing in the more abstracted representation. On the other hand, the economy provided by ignoring details of communication can allow one to increase detail elsewhere, and this may be a useful tradeoff. In this particular example, it is possible to increase the number of different types of faults that are considered in the most abstracted representation, and this allows the fault tolerance of the algorithm to be analyzed in greater detail (and reveals a bug in a published algorithm) (Lincoln and Rushby 1993).

As this example makes clear, abstraction is closely related to the modeling activity that is inherent in formal methods. The whole basis of formal methods is to create mathematical models of certain physical and computational phenomena and to make predictions about these phenomena through analysis of the models. Abstraction is concerned with how much, and what, detail to include in the model, and how to represent it. Validity of the predictions made through use of formal methods requires that the abstraction retains all salient details and that their formal representation is faithful to reality. Tractability of formal analysis, on the other hand, requires that the abstraction is ruthless in expelling all irrelevant detail. Ability to resolve this tension between too much and too little abstraction is the most important, and rarest, of the skills required to make effective use of formal methods.

3.2 THE VARIETIES OF FORMAL SPECIFICATIONS.

Formal methods embrace a variety of approaches that differ considerably in techniques, goals, claims, and philosophy. The previous section discussed some of the important differences, such as whether formal methods are used primarily for descriptive or for analytic purposes, the level of formality employed, and the stage(s) of the software development life cycle to which formal methods are applied. The different approaches to formal methods tend to be associated with different kinds of specification languages. Conversely, it is important to recognize that different specification languages are often intended for very different purposes and therefore cannot be compared directly to one another. Failure to appreciate this point is a source of much

misunderstanding. In the following sections some of the main varieties of specification languages are briefly introduced and their applications indicated.

3.2.1 Model-Oriented Specifications.

If specification or annotation of programs is the goal, then the formal notation employed should generally be close to, though more abstract than, that of programming, with operations changing values stored in an implicit system “state,” with data structures described fairly concretely, and with control described in operational terms. Formal notations with these characteristics are often described as model oriented, meaning that desired properties or behaviors are specified by giving a mathematical model that has those properties. For data structures, these models are often constructed from the notions of set theory using sets, functions, relations, etc. A pushdown stack, for example, can be modeled by a pair consisting of a natural number (the pointer) and a function (the stack) from natural numbers to the type of value being stacked. (This can be thought of as an array with the contents of the pushdown stack occupying positions $1 \dots \text{pointer}$; the empty stack is indicated by $\text{pointer} = 0$.) The “top” of the stack is the value of the function at the argument indicated by the pointer; the stack is “popped” by decrementing the pointer, and a value x is “pushed” on to the stack by incrementing the pointer and modifying the function so that it takes the value x at the argument indicated by the pointer.

NOTE: Representative notation for a model-oriented specification of a pushdown stack is the following, where a prime indicates the value of a variable in the new state following the operation and an unprimed name indicates the value in the old state prior to the operation. The \oplus operator indicates function modification (also called “overriding”).

<i>top</i>	is	<i>stack</i> (<i>pointer</i>), provided <i>pointer</i> > 0
<i>pop</i>	is	<i>pointer</i> ' = <i>pointer</i> - 1; provided <i>pointer</i> > 0
<i>push</i>	is	<i>pointer</i> ' = <i>pointer</i> + 1 \wedge <i>stack</i> ' = <i>stack</i> \oplus { <i>pointer</i> ' \mapsto χ }

To describe control, model-oriented notations for sequential programs generally provide sequential composition and if-then-else selection. Explicit loop constructs are not needed since their effects can usually be specified more abstractly using quantification. For example, in programming where a loop would be used to search for the least value stored in a pushdown stack, it can be formally specified that this value is the l such that (a) for all (\forall) natural numbers up to the pointer, the value of the stack at that point is no less than l , and (b) there exists (\exists) a natural number less than or equal to the pointer such that the value of the stack at that point equals l . Typical notation for specifying this is the following:

$$(\forall p: 1 \leq p \leq \text{pointer} \bullet \text{stack}(p) \geq l) \wedge (\exists p: 1 \leq p \leq \text{pointer} \bullet \text{stack}(p) = l).$$

A disadvantage of model-oriented specifications is that they can be overly prescriptive: suggesting how something is to be implemented, rather than just the properties it is required to have. For example, even though the specification of the least function does not prescribe an algorithm, it is

stated in terms of the pointer and array model, and so it would be fairly difficult to use this specification to establish correctness of an implementation that used linked lists instead.

3.2.2 Property-Oriented Specifications.

Preferred for program-level descriptions, specifications of early-life cycle products, such as requirements, commonly use property-oriented notations. These notations use an axiomatic style to state properties and relationships that are required to hold of the component being described, without suggesting how it is to be achieved. To specify a pushdown stack, for example, a property-oriented notation would state the relationships that are required to hold among the operations *top*, *pop*, and *push*: namely, that a *push* followed by a *pop* leaves the stack unchanged, and a *top* following a *push* returns the value that was pushed onto the stack.

Sequential control in property-oriented specifications is generally modeled by functional composition. For example, a *push* followed by a *pop* is specified by *pop (push [s; x])* rather than by the more operational *push (x); pop* (where the semicolon indicates sequential composition and the state of the stack is implicit in the operational specification). Iteration is generally modeled by quantification or recursion. For example, the least element in a stack can be specified by the following recursive function:

$$\textit{least}(s) = \textit{if empty}(s) \textit{ then } \infty \textit{ else } \textit{min}(\textit{top}(s), \textit{least}(\textit{pop}(s))) \textit{ endif.}$$

The treatment here of the empty stack (setting its least value to infinity) is a little suspect, though it can be made rigorous. Also, it has not been explained how the empty predicate is defined. Taking care of these difficulties in a fully satisfactory manner would require more space than is reasonable here, but they give a hint of the technical details that must be dealt with in a formal specification notation.

NOTE: Representative notation for a property-oriented specification of a pushdown stack is the following. Note that the value of the stack is supplied as an explicit argument (here *s*) to the operations, rather than being the implicit value of a program state.

<i>axiom pop_push</i>	is	<i>pop (push (s; x)) = s</i>
<i>axiom top_push</i>	is	<i>top (push (s; x)) = x</i>

Notice that although the stack operations are specified in a property-oriented style, this specification of the least function has an algorithmic flavor; the presentation in the previous section used a model-oriented specification for the data structure, but the specification of the least function was property-oriented (relative to the model of the data structure) rather than algorithmic. These mixed modes of expression are not uncommon.

An advantage of property-oriented over model-oriented specifications is the possibility to merely constrain certain relationships or values, without having to define them exactly. On the other hand, it is very easy to write conflicting constraints that cause the specification to become inconsistent. Inconsistent specifications are unimplementable, and are very dangerous because they can be used

to prove anything. Some Level 3 specification languages provide ways to ensure or demonstrate that property-oriented specifications are consistent.

3.2.3 Specifications for Concurrent Systems.

Concurrent and distributed systems can be specified in a variety of styles. One style takes some form of communication as primitive and has programming-like features for sending and receiving values. This style has a model-oriented flavor and is often referred to as process algebra. Another style takes shared variables as the primitive means of communication and often uses temporal logic to allow specification that a property should hold henceforth or eventually on some or all execution paths. This style has a property-oriented flavor. Methods associated with a kind of analysis known as model checking use one method (a kind of state machine) to specify the system concerned and another (a kind of temporal logic) to specify the properties required of it.

Further distinctions concern whether concurrent activities are considered to occur simultaneously (true concurrency) or alternately (interleaving concurrency), and whether consideration of time is restricted to the order in which events happen, or whether duration is considered (real-time logics).

3.3 THE VARIETIES OF FORMAL ANALYSIS.

Earlier, it was observed that one of the most significant differences among formal methods concerns whether their primary purpose is description or analysis. In fact, this distinction is too coarse: the kind of analysis must be known. The strongest kind of analysis is one that takes a formal description and predicts the behavior of a system satisfying that description. This is the kind of analysis that most closely corresponds to the use of applied mathematics and calculation in traditional engineering fields. If this is the goal desired of formal methods, then the modeling and notational techniques employed should favor efficient deduction, whereas more weight should be given to the readability of the notation when descriptive purposes are paramount. In an ideal world, one technique would serve both ends but, in the present state of the art, those notations that are considered most readable are much less tractable for automated reasoning than notations designed for that purpose. Conversely, notations designed for automated reasoning tend to have a rather austere and forbidding appearance. (The influence of notation can be illustrated by comparing Arabic with Roman numerals: for small numbers, at least, Roman numerals are more readable (for example, III is more suggestive of the concept *three* than is 3), but they are much less effective for calculation than are Arabic numerals.)

In between the purely descriptive uses of formal specifications and those that use automated deduction to make general predictions of behavior, there are many intermediate kinds of analysis that perform formal calculations to establish limited properties of a specification.

3.3.1 Consistency Analysis and Typechecking.

One kind of analysis does not attempt to deduce specific properties of the system described by a formal specification; instead it attempts to deduce whether such a system could exist. In other words, it checks whether the specification is sufficiently well formed to be a description of something. One very important well-formedness property is consistency: if a specification states

two contradictory things, then it cannot describe a real system and is therefore useless as a specification. Certain types of specifications lend themselves to systematic checks for consistency but, as always, there are tradeoffs involved. For example, specifications that allow new concepts to be introduced only by definition in terms of existing concepts can easily be checked for consistency; however, such specifications are purely constructive (i.e., strongly model oriented) and are unattractive for some purposes. In particular, constructive definitions are an unnatural way to state assumptions about the environment in which a system is to operate; axioms are more natural for this purpose, since the goal is to describe the environment, not to implement it. Consistency for axiomatic specifications can be established by showing that the axioms are true of some constructively defined implementation. This implementation need not be efficient or realistic, it just has to exist.

It is shockingly easy to write formal specifications that are inconsistent; consequently, any specifications offered in support of certification for safety-critical systems should be furnished with evidence for their consistency.

Some specification languages allow “types” to be given for entities appearing in specifications. Types are familiar from programming languages such as Ada or Pascal, where variables can be declared as integer or boolean and the compiler will generate an error message if an attempt is made to multiply an integer by a boolean. The types in a specification language can be more sophisticated (since they do not have to have a direct implementation), and the checks that are performed can be more elaborate. A computer program that checks specifications to ensure that entities are always used in ways compatible with their types is called a typechecker; it can be seen as a tool that performs a specialized kind of formal deduction (it attempts to prove the theorem “this specification is type-correct”). If the typechecker is allowed to use general-purpose theorem proving, rather than just perform algorithmic checks like a programming-language compiler, then the type system can become very sophisticated, and typechecking becomes a very powerful way to detect errors in a specification. Notably, higher-order logic has to use types to keep the logic consistent; types are technically optional for other kinds of logics (where they are sometimes called sorts). Ordinary set theory is untyped; when types are added (as in Z, for example), the result is a little awkward in that either some of the conveniences of set theory (e.g., nonhomogeneous sets) must be given up, or typechecking must be rather weak. Within these constraints, selection of a typed or an untyped specification language is often considered a matter of personal preference: some people value the early error detection of typechecking, others find the restrictions imposed by types to be irksome. However, those who prefer to forsake types should be expected to provide other, equally strong, evidence for the properties that would be established by typechecking.

Strong typechecking (the stronger the better) should always be required for formal specifications offered in support of certification for safety-critical systems.

3.3.2 Validating Formal Specifications.

Predictions are based on a mathematical model of the system; if the model is inaccurate, the predictions may not be true of the real system. It is therefore necessary to validate the accuracy of the model very carefully. Formal methods are no different in this regard than the mathematical

methods used in any other engineering field. Consistency analysis and typechecking provide evidence that a formal specification means something, but additional evidence is required to establish that it means what is intended. Reviews are one way to develop this evidence, but formal specifications can also support more analytical methods.

3.3.2.1 Animation.

One way to gain confidence in a formal specification is to test it on a few small examples. This kind of examination of a specification is sometimes called animation. Some kinds of formal specification can be executed using highly efficient forms of deduction, so that test cases can be run directly against the specification. Model-oriented specifications tend to lend themselves more naturally to direct execution than do property-oriented specifications. Note that executability may be at odds with other desirable properties of a specification (such as abstractness and nonprescriptiveness) (Hayes and Jones 1989). For specifications that cannot be executed directly, it may be desirable to construct a simulator or rapid prototype for testing purposes.

3.3.2.2 Formal Challenges.

Formal specifications can also be explored by posing and proving putative theorems called challenges: “if this specification says what it should, then the following ought to follow.” For example, suppose the operation of sorting a sequence had been specified; it might be asked whether sorting an already sorted sequence leaves the sequence unchanged (i.e., whether sorting is idempotent). That is, it might be asked whether:

$$\text{sort}(\text{sort}(x)) = \text{sort}(x)$$

is a theorem of the specification (assuming “sort” is a function that takes a sequence as argument and returns the sorted sequence as its value). Gerhart and Yelowitz (1976) describe how early formal specifications of sorting were deficient in that they required the output of the operation to be ordered but neglected to stipulate that it should also be a permutation of the input. An attempt to prove the theorem above would reveal such inadequacies.

Animation could examine the putative theorem for a few representative values of the input x , but a formal challenge would force consideration of all inputs. For some specifications of sorting, this could lead to the discovery that the putative theorem is unprovable; further examination might then lead to the notion of a stable sort (one that does not reorder elements of the sequence that are equivalent with respect to the ordering criterion, but distinguishable in other ways). It could then be decided whether stability was important to the application and, if so, could adjoin it as an additional requirement of the sorting specification.

Notice how this process of subjecting specifications to formal challenges probes the completeness as well as the correctness of specifications. Data from the Jet Propulsion Laboratory indicate that two-thirds of the defects in requirements specifications are omissions (Kelly, Sherif, and Hops 1992), so that systematic methods of exploring completeness are highly desirable.

Challenges can be undertaken at any level of formality, but it is beneficial for all those who write or review formal specifications to have experience in challenging specifications at Level 3 using a mechanized proof checker or theorem prover. Those who learn a formal specification language from textbooks or training courses, but who do not experiment with mechanically checked challenges, are in a position similar to that of learning a programming language without the opportunity to execute programs. In fact, their position is worse because experience with other programming languages is likely to help them learn a new one; whereas many of those learning a formal specification language are receiving their first exposure to formal methods, as well as to abstract and axiomatic forms of expression. Just as the failure of an obviously correct program teaches one that programming is difficult, so the discovery through dialog with a theorem prover that an expected property is not entailed by an obviously correct formal specification teaches one that specification is no easier than programming. This has to be learned individually, as only the personal shock of discovering egregious errors in personal specifications teaches one the requisite humility.

3.3.3 Predicting Behavior and Verifying Refinement.

Formal challenges probe a formal specification by asking whether it entails certain expected properties. Generally, these properties are special cases or fragments of the overall requirements. Once a specification has been sufficiently validated in this way, it is possible to examine it for the properties of real interest and to verify some steps in its refinement toward implementation.

The behavior of a system generally has many aspects and formal methods are usually not used to examine every aspect, but only those that are important to a particular analysis. For example, it may want to be known whether a system can deadlock or whether it can survive any single fault or whether a response is always delivered within a certain time. Some formal methods are especially well suited to the calculations necessary to predict certain kinds of properties; others are well suited to certain kinds of systems.

3.3.3.1 State Exploration.

Certain kinds of formal methods allow automatic, brute force exploration of all possible behaviors, provided there are not too many of them (the maximum number depends on many factors and increases as technology improves, but the typical range is from tens of thousands to tens of millions). Hardware and distributed algorithms such as protocols are particularly suitable for this kind of examination through exhaustive state exploration (Dill et al. 1992). Related technologies are known as model checking (Clarke, Grumberg, and Long 1994) and language inclusion (Har'El and Kurshan 1990).

A specification may admit too many behaviors for state exploration to succeed, but it may be possible to develop a downscaled version that can be examined in this way. For example, a communications protocol may be designed to move arbitrary data reliably over a faulty channel using sequence numbers that cycle through the range 0...255. For state exploration, the protocol can be downscaled to consider just one or two different data values and with sequence numbers restricted to 0 and 1. Experience indicates that examining all behaviors of a downscaled design is

often a more potent validation and assurance technique than examining some of the behaviors of the full design.

3.3.3.2 Verifying Desired Properties.

Formal methods provide the most searching examination and the strongest assurance when proofs are used to verify significant application properties. The process is essentially similar to that used in challenging specifications, except that the properties verified are of external significance, and the proofs are usually more difficult and longer. The basic idea is a formal specification of the requirement, design, or algorithm concerned, and also a formal statement of the property it is desired to satisfy is constructed, and then try to prove that the one implies the other. In practice, the first proof attempt seldom succeeds; instead it usually reveals the need for adjustments to the specification or to the statement of the desired property or the assumptions under which it is expected to hold. Generally, this process of adjustment is iterated several times as renewed proof attempts reveal additional subtleties. Experience shows that the process is always very enlightening, particularly when conducted with the full rigor of Level 3 formal methods, and leads to greatly enhanced understanding of the specifications and properties under examination.

If all goes well, the adjustments will converge and a satisfactory proof will finally be obtained. Subject to caveats on the fidelity of the modeling employed (these were mentioned at the end of section 2.2.1 and must be ensured by reviews and by validation as explained in the previous section) and on the proof being performed without error (this is where mechanized proof checking is valuable, but the mechanization is intended to enhance, not replace, human judgment and responsibility) the proof provides strong assurance that the specified artifact indeed satisfies the desired property. In addition, the process of formal modeling and proof generally provides other, incidental benefits: the discovery and correction of faults, complete enumeration of assumptions, sharpened statements of properties assumed or satisfied, streamlined arguments, and an enhanced understanding that can lead to improvements in design or assurance. Furthermore, formal specifications and proofs are a reusable intellectual resource that can be particularly valuable potentially, a corporate asset when design changes on the larger scale require modifications to the component under consideration: highly automated Level 3 verifications, in particular, can often verify slightly modified designs or properties with very little extra effort, and with the same degree of rigor as the original case—something that is very difficult to achieve with reviews.

3.3.3.3 Verifying Design Refinement.

Verification of requirements, designs, and algorithms against desired properties is typically an activity of the early life cycle. In the later life cycle, the task of verification is generally to demonstrate that the specification of a design at one level is implemented correctly by another at a more detailed level. Generally an abstraction function is used to translate the terms of the lower-level specification into those of the upper-level and the task is then to prove that this function has the properties of a homomorphism. The task becomes more complicated when the two levels operate at different granularities of time and especially when the timing relationship between them is variable (as, for example, in a pipelined microprocessor, where the implementation may take a variable number of cycles to complete an instruction, depending on whether the pipeline stalls and other details that are hidden at the upper level).

NOTE: The effectiveness of mechanized proof checkers and theorem provers for formal methods is advancing very rapidly. For example, mechanized verification of the microcode for a simple processor called "Tamarack" represented a significant challenge just five years ago (Joyce 1989), whereas it can now be done completely automatically in about five minutes (Cyrluk et al. 1994). Progress is uneven however, and the amount of human time and effort required to undertake a Level 3 analysis can vary by two orders of magnitude or more from one verification system to another. Potential users should be skeptical of the developer's own assessments of where their verification system stands in the power rankings, and of impressions gained from small examples; instead, they should evaluate candidate systems on full-size examples representative of the intended application.

Design refinement can be verified using any level of formality, but once the specifications get reasonably large, it is difficult to construct all the proof obligations to establish the homomorphism without mechanized assistance. On the other hand, the required proofs can be rather repetitive, so that investment in automation is often very worthwhile.

4. FORMAL METHODS AND CERTIFICATION.

This section is concerned with the use of formal methods in support of certification for airborne systems. First, the author's general recommendations are presented. These are followed by more specific recommendations in the context of an examination of the paragraphs concerning formal methods in RTCA/DO-178B.

4.1 GENERAL RECOMMENDATIONS.

Formal methods should be part of the education of every computer scientist and software engineer, just as the appropriate branch of applied mathematics is a necessary part of the education of all other engineers. Formal methods provide the intellectual underpinnings of this field; they can shape thinking habits and help direct approaches to problems along productive paths; they provide notations for documenting requirements and designs, and for communicating thoughts to others in a precise and perspicuous manner; and they provide one with analytical tools for calculating the properties and consequences of the requirements and designs that are documented.

However, it will be many years before even a small proportion of those working in industry have been exposed to a thorough grounding in formal methods, and it is simply impractical to demand large scale application of formal methods in airborne software and unnecessary too, since industry seems to be doing a mostly satisfactory job using nonformal methods. (The appalling safety record of the Airbus A320 aircraft [Boeing Commercial Airplane Group 1994] seems attributable to poor human factors rather than to specific software faults [Mellor 1994]).

Nonetheless, the author recommends that industry should be strongly encouraged to develop and apply formal methods that will permit more complete analysis and exploration of those aspects of design that are least well covered by present techniques. These arise in redundancy management, partitioning, and the synchronization and coordination of distributed components and primarily

concern fault tolerance, timing, concurrency, and nondeterminism. Scrupulous informal reviews, massive simulation, near-complete unit testing of components, and extensive all-up testing do not provide the same level of assurance in these cases as they do for sequential programs because the properties of interest are not manifest in individual components and because distributed execution is susceptible to subtle variations in timing and fault status that are virtually impossible to set up and cover in tests.

These formal analyses should be additional to those presently undertaken and can increase assurance without necessarily being totally comprehensive: the value of formal methods lies not in eliminating doubt but in circumscribing it. For example, in addition to all the other assurance techniques that may be applied, it will be valuable to prove that mode-switching logic does not contain states with no escape and that sensor data are distributed consistently, despite the presence of faults. These are not the only properties required of mode switching and sensor distribution, but they are among the most crucial and the most difficult to assure using traditional methods. To deal with such problems using current technology for formal methods it will often be necessary to abstract away irrelevant detail and possibly to simplify even relevant detail. Doing so while continuing to model the issues of real concern in a faithful way requires considerable talent and training. On the other hand, since only relatively small, albeit crucial, elements of the system will be subjected to analysis, the number of people required to possess that talent and training in formal methods will be small.

The benefit provided by these formal analyses is a complete exploration of a model of possible behaviors. Subject to the fidelity of the modeling employed (which must be established by extensive challenge and review) it will be assured that certain kinds of faults are not present at the level of description and stage of the life cycle considered. One source of doubt will have been eliminated and others posed more sharply. Admittedly, this does not guarantee that the implementation will not reintroduce the very faults that have been excluded by the formal analysis, but current practices seem effective at tracing implementations. As resources and capability permit, it will be worth seeing if formal methods can increase assurance for these aspects also, but initially cases where current practice seems weakest should be focused on, not where it seems effective. By that measure, other promising applications for formal methods are in the general area of requirements specification and analysis where current processes, though fairly effective, are ad-hoc and unstructured.

4.2 INTERPRETATION FOR RTCA/DO-178B.

The Requirements and Technical Concepts for Aviation, Inc. (RTCA) document known in the USA as RTCA/DO-178B and in Europe as EUROCAE ED-12B provides industry-accepted guidelines for meeting certification requirements for software used in airborne systems and equipment and is incorporated by reference into European and United States regulatory and advisory documents. RTCA/DO-178B provides guidelines and does not lay down specific certification requirements: those are based on existing regulations or special conditions decided by the certification authority in consultation with the applicant.

Formal methods are not specifically endorsed by RTCA/DO-178B (in contrast to certain other guidelines and standards that do recommend or require their use such as Interim Defence Standard 00-55: The Procurement of Safety Critical Software in Defence Equipment 1991) but are included among the alternative methods discussed in section 12.3 of RTCA/DO-178B.

Section 12.3 of RTCA/DO-178B on Alternative Methods states:

Some methods were not discussed in the previous sections of this document because of inadequate maturity at the time this document was written or because of limited applicability for airborne software. It is not the intention of this document to restrict the implementation of any current or future methods. Any single alternative method discussed in this subsection is not considered an alternative to the set of methods recommended by this document but may be used in satisfying one or more of the objectives of this document.

Also, section 12.3.1 on formal methods continues:

Formal methods involve the use of formal logic, discrete mathematics, and computer-readable languages to improve the specification and verification of software. These methods could produce an implementation whose operational behavior is known with confidence to be within a defined domain. In their most thorough application, formal methods could be equivalent to exhaustive analysis of a system with respect to its requirements. Such analysis could provide:

- Evidence that the system is complete and correct with respect to its requirements.
- Determination of which code, software requirements, or software architecture satisfy the next higher level of software requirements.

The goal of applying formal methods is to prevent and eliminate requirements, design, and code errors throughout the software development processes. Thus, formal methods are complementary to testing. Testing shows that functional requirements are satisfied and errors are detected, and formal methods could be used to increase confidence that anomalous behavior will not occur (for inputs that are out of range) or is unlikely to occur.

(Testing provides assurance that the normal behavior of the software is satisfactory; formal methods can extend that assurance by considering all possible behaviors, including those induced by rare or anomalous combinations of inputs and other circumstances.)

Section 12.3.1 of RTCA/DO-178B recognizes different levels of formality and rigor in applications of formal methods. (These levels do not quite correspond to those introduced earlier as Levels 1, 2, and 3. However, the exact demarcation of levels is unimportant; the main point is that formal

methods can be applied at many levels of rigor, and it is important to consciously select one that is appropriate for the intended application).

Formal methods include these increasingly rigorous levels:

- Formal specification with no proofs
- Formal specifications with manual proofs
- Formal specifications with automatically checked or generated proofs

The use of formal specifications alone forces requirements to be unambiguous. Manual proof is a well-understood process that can be used when there is little detail. Automatically checked or generated proofs can aid the human proof process and offer a higher degree of dependability, especially for more complicated proofs.

Section 12.3 of RTCA/DO-178B provides guidance in using an alternative method. It states:

An alternative method cannot be considered in isolation from the suite of software development processes. The effort for obtaining certification credit of an alternative method is dependent on the software level and the impact of the alternative method on the software life cycle processes. Guidance for using an alternative method includes:

- a. An alternative method should be shown to satisfy the objectives of this document.
- b. The applicant should specify in the Plan for Software Aspects of Certification and obtain agreement from the certification authority for
 1. the impact of the proposed method on the software development processes. (Software development processes, discussed in section 5 of RTCA/DO-178B, include software requirements analysis, design, coding, and integration.)
 2. the impact of the proposed method on the software life cycle data. (Software life cycle data, which are discussed in section 11 of RTCA/DO-178B, are produced “to plan, direct, explain, define, record, and provide evidence of activities throughout the software life cycle”.)
 3. the rationale for use of the alternative method that shows that the system safety objectives are satisfied. The rationale should be substantiated by software plans, processes, expected results, and evidence of the use of the method.

The effort required to satisfy these guidance items will depend on the extent to which formal methods replace, rather than merely supplement, traditional life cycle processes and data. Note,

however, that even if formal methods are truly supplementary to the traditional processes, there will still be some impact on the life cycle processes and data that the applicant should discuss and explain. For example, if descriptive formal methods are used to supplement traditional documentation, then the questions have to be asked: which is the primary description, how is consistency established between them (and maintained through changes), and which is the target of the reviews and analyses performed for verification?

The guidance for alternative methods in section 12.3 of RTCA/DO-178B does not explicitly call for an assessment of the impact of the proposed method on verification activities. However, this is implicit in the requirement (quoted in paragraph 12.3.b[2]) to consider the impact of the method on life cycle data since these data record the results of verification activities. Verification is one of what RTCA/DO-178B calls the integral processes, which supports the development processes by ensuring the correctness and quality of all processes and the delivered software.

The integral processes comprise software verification, software configuration management, software quality assurance, and certification liaison. They are discussed in sections 6 through 9 of RTCA/DO-178B; in particular, technical guidance on verification is found in section 6. (The guidance on software quality assurance found in section 8 is chiefly concerned with monitoring and recording the processes recommended in the other sections.)

5. SOFTWARE VERIFICATION PROCESS.

RTCA/DO-178B states:

Verification is not simply testing. Testing, in general, cannot show the absence of errors. As a result, the following subsections use the term verify instead of test when the software verification objectives being discussed are typically a combination of reviews, analyses, and tests.

A distinction between reviews and analyses is pertinent when considering formal methods. On the topic of Software Reviews and Analyses RTCA/DO-178B states:

Reviews and analyses are applied to the results of the software development processes and software verification process. One distinction between reviews and analyses is that analyses provide repeatable evidence of correctness and reviews provide a qualitative assessment of correctness.

(A draft version of RTCA/DO-178B said “The primary distinction between reviews and analyses is that analyses provide repeatable evidence, and reviews provide a group consensus of correctness.”)

The significant attribute of formal methods is that they allow certain questions to be settled by calculation, that is by analysis, which informal methods must resolve by means of reviews.

An applicant who offers descriptive formal methods might argue that formal analysis is not his goal and that reviews provide an adequate means of verification for his purpose. This could be

acceptable if formal methods are offered only as a supplement to traditional documentation and if the traditional verification processes are applied to that documentation (though it would then be natural to question the purpose of offering formal methods); it could also be acceptable if formal methods are used for limited purposes, such as describing data structures. One should be skeptical, however, of the reliability of reviews when they are applied to substantial formal specifications involving, for example, significant numbers of axioms or operations specified by complex pre- and post-conditions or constructions with subtle semantics (e.g., schemas in the language Z). The problem is that it is difficult to provide objective evidence that the authors of a specification can reliably express themselves in such forms and that its reviewers can interpret them correctly. The problem is compounded by the fact that such specifications often contain technical errors (the equivalent of coding bugs) that can render them inconsistent or meaningless.

It is the author's opinion that an essential step in ensuring an effective review process for formal specifications is to require that they are subjected to stringent (and preferably mechanized) analysis before they are submitted to reviews. The purpose of the analysis is to eliminate as large a class of potential faults as possible by purely formal means (i.e., by calculational processes), so that the review process may concentrate on the intellectual substance of the specification. Some specific forms of analysis that should be considered (in ascending order of stringency) are

- parsing.
- typechecking (there are many degrees of stringency possible here; the most stringent generally require use of theorem proving).
- well-formedness checking for definitions (i.e., assurance of conservative extension).
- demonstration of consistency for axiomatic specifications (i.e., exhibition of models).
- animation (i.e., construction of an executable prototype from the formal specification so that it can be subjected to experiment). This form of analysis has a rather different character than the others listed here and should be used for specific purposes that are defined beforehand otherwise it can degenerate into hacking.
- state exploration (i.e., exploring all the behaviors of a possibly downscaled version of the specification).
- formal challenges (i.e., posing and proving putative theorems that should be entailed by the specification).

Mechanically supported analyses of the kinds suggested above are potent forms of fault detection for formal specifications. In many projects it will also be worthwhile to develop additional forms of mechanized analysis to check for specific classes of faults. By these means, it can be ensured that the formal specifications submitted for review are free of gross defects and the reviewers can focus on deeper issues.

The question remains: how much confidence can one have in reviews of formal specifications by personnel who may not be experts in formal methods? Trust in the integrity of the review process seems the best way to settle this. Currently, reviews are conducted using checklists with items such as “do you consider the requirements are complete?” , and it will be necessary to add items such as “do you consider that you have been able to fully comprehend the formal specification?” The assurance that participants do fully comprehend a formal specification may be enhanced if the suggestions of Parnas and Weiss (1985) are followed. For example, someone other than the author of a specification should be expected to explain it during the review, and the author should pose questions to the reviewers (rather than vice versa).

The rationale submitted to satisfy section 12.3.b(3) of RTCA/DO-178B should clearly state the analyses that are required to be completed prior to reviews and should describe the class of faults that are detected by means of these analyses and whether the detection is certain or merely likely. The number and stringency of the analyses performed may be determined by the criticality and sophistication of the formal specifications considered. The rationale should also provide evidence that the applicant’s process for reviewing formal specifications is effective.

Although descriptive formal methods have value, it is when formal methods exploit the power of calculational, and especially automated, forms of analysis that their singular potential is best realized. For example, section 6.3.1 (Reviews and Analyses of High-Level Requirements), RTCA/DO-178B, requires:

Accuracy and Consistency: The objective is to ensure that each high-level requirement is accurate, unambiguous, and sufficiently detailed and that the requirements do not conflict with each other. (Similar considerations apply to lower-level requirements described in section 6.3.2 of RTCA/DO-178B).

Some aspects of consistency and nonconflict can be established mechanically for formal specifications by strong typechecking and related analyses.

At a more detailed level, the verification objectives stated in section 6.3.1 of RTCA/DO-178B include:

- Algorithm Aspects - The objective is to ensure the accuracy and behavior of the proposed algorithms, especially in the area of discontinuities.

Suitable verification objectives for analytic formal methods would be to demonstrate that certain fundamental algorithms and architectural mechanisms involving complex behavior (i.e., behavior with large numbers of discontinuities) satisfy their corresponding high-level requirements. Other suitable verification objectives for strong kinds of formal analyses may be to discharge the requirements of RTCA/DO-178B, section 6.2.d, which states:

When it is not possible to verify specific software requirements by exercising the software in a realistic test environment, other means

should be provided and their justification for satisfying the software verification process objectives should be defined in the Software Verification Plan.

Circumstances, “when it is not possible to verify specific software requirements by exercising the software in a realistic test environment,” are likely to include those where complex interactions produce very large numbers of possible behaviors (e.g., in coordination of real-time or concurrent processes and in redundancy management), and analytic formal methods may be the most effective other means of verification in these cases. RTCA/DO-178B seems to agree that formal methods are well suited to such complex behaviors, and in section 12.3.1 states

...that formal methods may be applied to software requirements that:

Are safety-related.

Can be defined by discrete mathematics.

Involve complex behavior, such as concurrency, distributed processing, redundancy management, and synchronization.

The specific benefit provided by formal methods is that they allow complex behaviors to be analyzed (by means of proofs or state exploration), rather than reviewed and analyzed in their totality and sampled as by testing or simulation. Thus, the benefit derives from formal analysis, not from formal specification alone. Formally specifying the individual state machines at either end of a protocol, for example, adds little to the understanding that is needed to calculate the combined behavior to ensure that they accomplish the desired goal.

The author’s recommendation is that those aspects of design that involve complex behavior should be provided with at least the level of formal description and analysis that would be found in a refereed computer science journal. That is, for the mechanisms of redundancy management, a specification of the relevant algorithms should be provided together with the fault assumptions, the fault masking or recovery objectives, and a proof that the algorithms satisfy the objectives subject to the assumptions. This level of rigor of presentation is what was earlier called a Level 1 formal method; it is less rigorous than any of the levels of formality mentioned in section 12.3.1 of RTCA/DO-178B. Nonetheless, this level of rigor would be a distinct improvement on current practice.

Beyond this modest step, the extent to which quality control and assurance might be further enhanced by increasing the level of formal rigor employed or increasing the number of stages of the life cycle subjected to formal analysis should be considered.

When the concern is to establish that certain tricky or crucial aspects of design are correctly handled by the algorithms and architectural mechanisms employed, there is little advantage to Level 2 formal methods over Level 1: it is the proofs that matter, and both levels employ the traditional kind (presented and checked informally by hand); all that Level 2 would add is a fixed syntax for

the specification language and possibly some built-in models of computation and concurrency. These last may be a mixed blessing: useful if they match the needs of the problems considered, otherwise an obstacle to be overcome.

However, if Level 2 formal methods add very little in this domain, Level 3 may add a great deal. The focus will be on difficult problems, where a large number of potential behaviors must be considered. This is why the applicant has decided to use formal methods and the proofs may be expected to be replete with boundary conditions and case analyses. These are precisely the kinds of arguments where informal reasoning may be expected to go astray. For example, the published proof for one synchronization algorithm (Lamport and Melliar-Smith 1985) has flaws in its main theorem and in four of its five lemmas (Rushby and von Henke 1993). The flaws in this example were discovered while undertaking formal analysis at Level 3 and suggest the benefits that may be derived from this level of rigor.

The value of undertaking mechanically checked proofs is that the dialog with a proof checker forces examination of all cases to the argument. On its own, a mechanically checked proof is not a means of compliance with a certification basis and concern that “the theorem prover has not itself been proved correct” is not an obstacle to deriving great benefit and additional assurance by applying Level 3 formal methods in this domain. The analysis produced through dialog with any adequately validated proof-checker will be considerably more complete and reliable (and repeatable) than one produced without such aid. It is the ultimate walk-through but the certificate that comes from a mechanized proof checker should not be accepted as unsupported evidence of fitness any more than should other computer-assisted calculations, such as those of aerodynamic properties or of mechanical stress. In the author’s opinion, the analysis developed with the aid of a theorem prover should also be rendered into a clear and compelling semiformal (i.e., Level 1) argument that is subjected to intense human review, and it is the combination of stringent mechanical and human scrutiny (and other evidence, such as tests) that should be considered in certification. (For this reason, the author does not endorse the requirement in UK Interim Defence Standard 00-55: The Procurement of Safety Critical Software in Defence Equipment (1991), paragraph 32.2.3 that a second mechanically checked proof using a “diverse tool” should be required. The resources required would be better expended on diverse analysis and on human scrutiny of the argument and modeling employed.)

The construction of a mechanically checked proof that certain algorithms and architectural mechanisms accomplish certain goals subject to certain assumptions addresses only part of the problem. It is also necessary to validate the modeling employed. That is, the applicant needs to provide evidence that the model of computation employed, and the statements of the assumptions made and of the goals to be achieved, are all true in the intended interpretation. It is also necessary to provide evidence that the algorithm and architectural mechanisms considered in the proof are correctly implemented. There is a tension between these concerns: it is generally easier to validate models that make a few broad and abstract assertions (e.g., it is possible for a nonfaulty processor to read the clock of another nonfaulty processor with at most a small error ϵ) than those that make many detailed ones (e.g., that talk about specific mechanisms for reading clocks and the behavior of particular interface registers), but the gap between the verified specification and its implementation will be greater in the former case.

Since the assurance objective of this analysis is to ensure that there are no conceptual flaws in the basic algorithms and mechanisms, the author's opinion is that credibility of validation should take precedence over proximity to implementation. This argues for performing the analysis early in the life cycle and using abstract modeling (i.e., suppressing all detail judged irrelevant). Validation should be accomplished by peer review supported by analyses that demonstrate, for example, that axiomatic specifications are consistent (i.e., have a model), that intended models are not excluded (e.g., that clocks that keep perfect time satisfy the axioms for a good clock), that definitions are well formed, and that expected properties (i.e., challenges) can be proven to follow from the specification. Concern that implementations are faithful to their verified specifications is a separate problem and can be handled using either formal methods or traditional techniques for verification and validation. In the author's opinion traditional techniques are likely to be adequate. The evidence seems to be that it is the basic mechanisms and algorithms that have been flawed, not their implementations.

6. CONCLUSION.

The recommendations presented above may seem modest to those who believe that formal methods should be used more extensively (for example, in the manner required by UK Interim Defence Standard 00-55). They may also seem a retreat from the traditional goals of formal verification: there would be no claim of proving correctness and no ambition to apply formal methods from top to bottom (i.e., from requirements down to code or gates). Rather, the goal would be to establish that certain properties hold, and certain conceptual faults are absent, in formal models of some of the basic mechanisms necessary to safe operation of the system. These may seem small claims, but they are the claims that the author considers are least well supported by current practice and that cause the most concern, since they are the most fundamental.

Those who argue that more should be required, that formal methods should be carried down to code or gates, or that formal specifications should be used as part of the software engineering process need to provide evidence that this will increase assurance in an industry that has an excellent record of accomplishment using traditional methods. They also need to provide evidence that resources expended on formal methods would not be better spent on other forms of assurance.

On the other hand, these recommendations may seem excessive to some readers: it is proposed that the most stringent kinds of formal methods should be applied to the hardest problems of design. These pose tough challenges, to be sure, but how could anything less challenging be expected to improve a process that is already very effective? Furthermore, notice that although these challenges are tough, they are relatively few in number and small in scale, and can therefore be undertaken by a small (though highly skilled) team of people. The tools that are currently available to support these ambitious applications of formal methods are not ideal but are adequate to the task.

7. BIBLIOGRAPHY.

- Statistical Summary of Commercial Jet Aircraft Accidents, Worldwide Operations, 1959-1993, Airplane Safety Engineering (B-210B), Boeing Commercial Airplane Group, Seattle, WA, March 1994.
- Butler, Ricky W. and George B. Finelli, "The Infeasibility of Experimental Quantification of Life-Critical Software Reliability," IEEE Transactions on Software Engineering, 19(1):3-12, January 1993.
- Biddle, Wayne, Barons of the Sky: From Early Flight to Strategic Warfare, the Story of the American Aerospace Industry, Simon and Schuster, New York, NY, 1991.
- Brazendale, J. and A. R. Jeffs, "Out of Control: Failures Involving Control Systems," High Integrity Systems, 1(1):67-72, 1994.
- Burns, R. W., "Genius at Work," IEE Review, 39(5):187-189, September 1993.
- Clarke, Edmund M. et al., "Verification of the Futurebus+ Cache Coherence Protocol," In David Agnew, Luc Claesen, and Raul Camposano, eds., CHDL '93: 11th Conference on Computer Hardware Description Languages and their Applications, pages 5-20, Ottawa, Canada, April 1993.
- Clarke, E., O. Grumberg, and D. Long, "Verification Tools for Finite State Concurrent Systems," in J. W. de Bakker, W. P. de Roever, and G. Rozenberg, eds., A Decade of Concurrency, pages 124-175, REX Workshop, Mook, The Netherlands, June 1994. Volume 803 of Lecture Notes in Computer Science, Springer-Verlag.
- Cooper, Henry S. F. Jr., The Evening Star: Venus Observed, Farrar Straus Giroux, New York, NY, 1993.
- Cyrluk, D., S. Rajan, N. Shankar, and M. K. Srivas, "Effective Theorem Proving for Hardware Verification," in Ramayya Kumar and Thomas Kropf, eds., Preliminary Proceedings of the Second Conference on Theorem Provers in Circuit Design, pages 287-305, Bad Herrenalb (Blackforest), Germany, September 1994.
- Dill, David L., Andreas J. Drexler, Alan J. Hu, and C. Han Yang, "Protocol Verification as a Hardware Design Aid," 1992 IEEE International Conference on Computer Design: VLSI in Computers and Processors, pages 522-525, IEEE Computer Society, Cambridge, MA, 1992.
- Dornheim, Michael A., "X-31 Flight Test to Explore Combat Agility to 70 Degrees AOA," Aviation Week and Space Technology, pages 38-41, March 1991.
- Eckhardt, Dave E., Alper K. Caglayan, John C. Knight, Larry D. Lee, David F. McAllister, Mladen A. Vouk, and John P. J. Kelly, "An Experimental Evaluation of Software Redundancy as a

- Strategy for Improving Reliability,” IEEE Transactions on Software Engineering, 17(7):692-702, July 1991.
- Hecht, Myron J., “Fault-Tolerant Software,” Digital Systems Validation Handbook-Volume II, Chapter 9, DOT/FAA/CT-88/10, Federal Aviation Administration Technical Center, Atlantic City, NJ, February 1989.
- Advisory Circular 25.1309-1A, System Design and Analysis, Federal Aviation Administration, June 21, 1988.
- Digital Systems Validation Handbook-Volume II, Federal Aviation Administration Technical Center, Atlantic City, NJ, DOT/FAA/CT-88/10, February 1989.
- Guiho G. and C. Hennebert, “SACEM Software Validation,” 12th International Conference on Software Engineering, IEEE Computer Society, pages 186-191, Nice, France, March 1990.
- Gerhart, S. L. and L. Yelowitz, “Observations of Fallibility in Modern Programming Methodologies,” IEEE Transactions on Software Engineering, SE-2(3):195-207, September 1976.
- Har’El, Zvi and Robert P. Kurshan, “Software for Analytical Development of Communications Protocols,” AT&T Technical Journal, 69(1):45-59, January/February 1990.
- Hayes, I. J. and C. B. Jones, “Specifications Are not (Necessarily) Executable,” IEE/BCS Software Engineering Journal, 4(6):320-338, November 1989.
- Hecht, Herbert, “Rare Conditions: An Important Cause of Failures,” COMPASS ‘93: Proceedings of the Eighth Annual Conference on Computer Assurance, pages 81-85, IEEE Washington Section, Gaithersburg, MD, June 1993.
- Interim Defence Standard 00-55: The Procurement of Safety Critical Software in Defence Equipment, UK Ministry of Defence, April 1991. Part 1, Issue 1: Requirements; Part 2, Issue 1: Guidance.
- Jones, Cliff B., “Systematic Software Development Using VDM,” Prentice Hall International Series in Computer Science, Prentice Hall, Hemel Hempstead, UK, second edition, 1990.
- Joyce, Jeffrey, Multi-Level Verification of Microprocessor-Based Systems, PhD Thesis, University of Cambridge, December 1989.
- Keutzer, Kurt, “The Need for Formal Verification in Hardware Design and what Formal Verification Has not Done for me Lately,” In Phillip Windley, ed., Proceedings of the 1991 International Workshop on the HOL Theorem Proving System and its Applications, IEEE Computer Society, pages 77-86, Davis, CA, August 1991.

- Knight, J. C. and N. G. Leveson, "An Experimental Evaluation of the Assumption of Independence in Multiversion Programming," IEEE Transactions on Software Engineering, SE-12(1):96-109, January 1986.
- Kasuda, Rick and Donna Sexton Packard, "Spacecraft Fault Tolerance: The Magellan Experience," In Robert D. Culp and George Bickley, eds., Proceedings of the Annual Rocky Mountain Guidance and Control Conference, pages 249-267, Keystone, CO, February 1993. Volume 81 of Advances in the Astronautical Sciences, American Astronautical Society.
- Kelly, John C., Joseph S. Sherif, and Jonathan Hops, "An Analysis of Defect Densities Found During Software Inspections," Journal of Systems Software, 17:111-117, 1992.
- Lamport, Leslie and Stephan Merz, "Specifying and Verifying Fault Tolerant Systems," In H. Langmaack, W.-P. de Roever, and J. Vytupil, eds., Formal Techniques in Real-Time and Fault-Tolerant Systems, pages 41-76, L'ubeck, Germany, September 1994. Volume 863 of Lecture Notes in Computer Science, Springer-Verlag.
- Lamport, L. and P. M. Melliar-Smith, "Synchronizing Clocks in the Presence of Faults," Journal of the ACM, 32(1):52-78, January 1985.
- Lamport, Leslie, Robert Shostak, and Marshall Pease, "The Byzantine Generals Problem," ACM Transactions on Programming Languages and Systems, 4(3):382-401, July 1982.
- Lincoln, Patrick and John Rushby, "A Formally Verified Algorithm for Interactive Consistency Under a Hybrid Fault Model," Fault Tolerant Computing Symposium 23, pages 402-411, IEEE Computer Society, Toulouse, France, June 1993.
- Lloyd, E. and W. Tye, Systematic Safety: Safety Assessment of Aircraft Systems, Civil Aviation Authority, London, England, 1982. Reprinted 1992.
- Leveson, Nancy G. and Clark S. Turner, "An Investigation of the Therac-25 Accidents," IEEE Computer, 26(7):18-41, July 1993.
- Lutz, Robyn R., "Analyzing Software Requirements Errors in Safety-Critical Embedded Systems," IEEE International Symposium on Requirements Engineering, pages 126-133, San Diego, CA, January 1993.
- Mackall, Dale A., "Development and Flight Test Experiences with a Flight-Crucial Digital Control System," NASA Technical Paper 2857, NASA Ames Research Center, Dryden Flight Research Facility, Edwards, CA, 1988.
- Mellor, Peter, "CAD: Computer-Aided Disaster," High Integrity Systems, 1(2):101-156, 1994.

- Miller, Steven P. and Mandayam Srivas, "Formal Verification of the AAMP5 Microprocessor: A Case Study in the Industrial Use of Formal Methods," WIFT '95: Workshop on Industrial-Strength Formal Specification Techniques, IEEE Computer Society, 1995.
- Owre, Sam, John Rushby, Natarajan Shankar, and Friedrich von Henke, "Formal Verification for Fault-Tolerant Architectures: Prolegomena to the Design of PVS," IEEE Transactions on Software Engineering, 21(2), February 1995.
- Parnas, David L. and David M. Weiss, "Active Design Reviews: Principles and Practices," 8th International Conference on Software Engineering, IEEE Computer Society, pages 132-136, London, UK, August 1985.
- Rumbaugh, James, Michael Blaha, William Premerlani, Frederick Eddy, and William Lorenzen, Object-Oriented Modeling and Design, Prentice Hall, Englewood Cliffs, NJ, 1991.
- RTCA/DO-178B, "Software Considerations in Airborne Systems and Equipment Certification," Requirements and Technical Concepts for Aviation, Washington, DC, December 1992.
- Rushby, John, "Formal Methods and Digital Systems Validation for Airborne Systems," Technical Report SRI-CSL-93-7, Computer Science Laboratory, SRI International, Menlo Park, CA, December 1993 (also as NASA Contractor Report 4551, December 1993).
- Rushby, John and Friedrich von Henke, "Formal Verification of Algorithms for Critical Systems," IEEE Transactions on Software Engineering, 19(1):13-23, January 1993.
- Spivey, J. M., ed., "The Z Notation: A Reference Manual," Prentice Hall International Series in Computer Science, Prentice Hall, Hemel Hempstead, UK, second edition, 1993.
- Vincenti, Walter G., "What Engineers Know and How They Know It: Analytical Studies from Aeronautical History," Johns Hopkins Studies in the History of Technology, The Johns Hopkins University Press, Baltimore, MD, 1990.

8. GLOSSARY.

These explanations are provided to help the nonspecialist. They are intended to reflect the technical uses of the terms considered but do not attempt to incorporate subtleties that concern the specialist.

ABSTRACTION. The process of simplifying certain details of a system description or model so that the main issues are exposed. Abstraction is the key to gaining intellectual mastery of any complex system and a prerequisite to effective use of formal methods. It requires great skill and experience to use abstraction to its best effect.

In formal methods, abstraction is part of the process of developing a mathematical model that is a simplification or approximation of reality but that retain the properties of interest. In physics, for example, it is customary to model a moving object as a point mass and to ignore its shape. Similarly in the case of a flight-control system, one can analyze properties of the clock synchronization algorithm or the redundancy management mechanisms by abstracting these away from the larger and more complex system in which they are embedded.

CORRECTNESS. The property that a system does what it is expected and required to do. Formal methods cannot establish correctness in this most general sense because they deal with formal models of the system that may be inaccurate or incomplete and with formal statements of requirements that may not capture all expectations. The difference between the real and modeled worlds is a potential source of error that attends all uses of mathematical modeling in engineering (e.g., in numerical aerodynamics or stress calculations) and that must be controlled by validating the models concerned. The difference between expectations and documented requirements is another problem that attends all engineering activities. Formal methods provide ways to make the specifications of assumptions and requirements precise; formal validation (q.v.) can then be used to ensure that the specifications are adequately complete and correct.

Correctness does not ensure safety of other critical properties, since the system requirements and expectations may not address these issues (correctly or completely). System requirements usually describe functional properties (i.e., what the system is to do); it is necessary to establish nonfunctional properties such as safety and security (which often describe what the system is not to do) by separate scrutiny (based, e.g., on hazard analysis or threat analysis). Formal methods can be used in these processes.

DESIGN FAULTS. Mistakes in the design of the system or in the understanding of its requirements and assumptions that cause it to do the wrong thing or fail in certain circumstances. Also called generic faults. Modular redundancy provides no protection against these faults.

FORMAL LOGIC. Symbolic notation equipped with rules for constructing formal proofs (q.v.). Formal logic consists of a language for writing statements and syntactic rules of inference for constructing proofs using these statements. Formal logic supports a form of reasoning that does not rely on the subjective interpretation of the symbols used. For example, “All \diamond s are $\#$; this \odot is a \diamond ; therefore, this \odot is $\#$ ” is sound reasoning, no matter what the symbols mean. Because

they do not depend on intuition, formal proofs can be developed or checked by computer (see theorem proving).

There are many formal logics; they differ in what concepts they can express, and in how difficult it is to discover or check proofs. Propositional logic, first-order logic, higher-order logic, the simple theory of types, and temporal logic are all examples of formal logics that find applications in formal methods. These logics are generally augmented with certain theories defined within them that provide definitions or axiomatizations for useful mathematical concepts such as sets, numbers, state machines, etc.

FORMAL PROOF AND FORMAL DEDUCTION. Formal deduction is the process of deriving a sentence expressed in a formal logic form other through application of one or more rules of inference. In the example above, formal deduction allows one to derive the sentence “this \odot is #” from the two sentences “All \diamond s are #” and “this \odot is a \diamond .”

A formal proof is a demonstration that a given sentence (the theorem) follows by formal deduction from given (i.e., assumed) sentences called premises.

FORMAL METHODS. Methods that use ideas and techniques from mathematical or formal logic (q.v.) to specify and reason about computational systems (both hardware and software).

FORMAL SPECIFICATION. A description of some computational system expresses in a notation based on formal logic. Generally, the specification states certain assumptions about the context in which the system is to operate (e.g., laws of physics, properties of subsystems and of systems with which the given system is to interact) and certain properties required of the system. A requirements specification need specify no more than this; a design specification will specify some elements of how the desired properties are to be achieved (e.g., algorithms and decomposition into subsystems).

FORMAL VALIDATION. A process for gaining confidence that top-level formal specifications of requirements and assumptions are correct. Formal verification (q.v.) cannot be applied at these levels because there are no higher-level requirements or more basic assumptions against which to verify them: processes of review and examination must be used instead. Formal validation consists of challenging the formal specifications by proposing and attempting to prove theorems that ought to follow from them (i.e., “if I’ve got this right, then this ought to follow”).

FORMAL VERIFICATION. The process of showing, by means of formal deduction, that a formal design specification satisfies its formal requirements specification. The formal description of a design and its assumptions supply the premises, and the requirements supply the theorem to be proved. In hierarchical developments, assumptions and designs at one level become requirements at another, so the formal verification process can be repeated through many levels of design and abstraction. At the topmost level, validation (q.v.) must be employed.

THEOREM PROVING AND PROOF CHECKING. Given a putative theorem and its premises, a theorem prover attempts to discover a proof that the theorem follows from the premises; on the

other hand, a proof checker simply checks that a proof is valid according to the rules of deduction for the logic concerned. Both these processes can be automated. A theorem prover is a computer program that uses search, heuristics, and user-supplied hints to guide its attempt to discover a proof. A proof checker is a computer program that is used interactively: a human user proposes proof steps and the proof checker checks they are valid and carries them out. The most effective automated assistance for formal methods is generally obtained by a hybrid combination of these approaches: the user proposes fairly big steps and the proof checker uses theorem proving techniques to fill in the gaps and take care of the details. Examples of theorem provers include Otter, Nqthm, PTTP, RRL, and TPS. Examples of proof checkers include Automath, Coq, HOL, Isabelle, and Nuprl. Hybrids include Eves, IMPS, PC-Nqthm, and PVS. Other forms of automated analysis that can be applied to formal specifications include model checking, language inclusion, and state exploration; examples of systems that perform these analysis are SMV, COSPAN, and MurØ.

INDEX

- abstraction, 1, 10, 11, 13, 19, 20, 27, 41, 43
- accuracy, 34
- alternative methods, 30, 32
- animation, 25, 33
- anomalous behavior, 18, 30
- assurance, 1, vi, 1, 3, 4, 6, 7, 8, 10, 11, 12, 17, 27, 29, 30, 32, 33, 34, 35, 36, 37
- axioms, 8, 14, 15, 24, 33, 37
- certainty of calculation, 9
- certification, 1, vi, 4, 5, 17, 24, 28, 29, 31, 32, 36
- certification credit, 31
- common mode failure, 2
- completeness, 25
- complexity, 1, 2, 4, 5
- consistency, 24, 32, 33, 34
- correctness, 1, 19, 22, 25, 32, 37, 41
- data structures, 12, 21, 33
- design fault, 1, 2, 3, 5, 12
- distributed algorithm, 20, 26
- documentation, 10, 11, 19, 32, 33
- fault injection, 7
- formal challenges, 25, 33
- formal deduction, 1, 8, 9, 24, 42, 43
- formal methods, 1, vi, 1, 6, 7, 8, 9, 10, 11, 12, 15, 16, 17, 18, 19, 20, 23, 26, 27, 28, 29, 30, 32, 33, 34, 35, 36, 37, 41, 42, 43
- formal specifications, 8, 23, 24, 25, 26, 27, 31, 33, 34, 37, 43
- formality, 12, 13, 15, 20, 26, 28, 30, 35
- hazard analysis, 4, 6, 42
- inductive assertion, 16
- levels of formality, 15, 30, 35
- mechanized analysis, 15, 33
- model, 1, 7, 9, 10, 15, 20, 21, 22, 23, 24, 26, 29, 36, 37, 41
- model checker, 15
- modeling, 6, 7, 9, 10, 18, 20, 23, 27, 29, 36, 37, 42
- modular redundancy, 2
- parsing, 33
- physical system, 2, 5, 6
- proof checker, 9, 15, 26, 28, 36, 43
- quality assurance, 3, 4, 32
- quality control, 3, 11, 17, 35
- rationale for formal methods, 1, vi, 1
- redundancy management, 3, 10, 17, 35, 41
- reliability, 5, 6, 7, 33
- RTCA/DO, iv, 1, 4, 6, 17, 18, 28, 29, 30, 31, 32, 34, 35, 41
- selected components, 12
- selected properties, 12
- selection, 1, 9, 10, 11, 21, 24
- sequential composition, 14, 15, 21, 22
- set theory, 21, 24
- software configuration management, 32
- software level, 4, 31
- software quality assurance, 32
- software requirements, 30, 31, 35
- software verification, 4, 32, 35
- specification notations, 16
- state exploration, 26, 33, 35
- testing, 4, 5, 8, 19, 25, 29, 30, 32, 35
- theorem prover, 9, 12, 15, 16, 26, 28, 36, 43
- theorem proving, 1, 19, 20, 24, 33, 42, 43
- typechecker, 24
- typechecking, 24, 25, 33, 34
- verification, 1, vi, 4, 15, 16, 17, 18, 19, 20, 27, 28, 30, 32, 33, 34, 35, 37, 43



