

IEEE SOFTWARE 0740-7459/97/\$10.00 © 1997 IEEE

Vol. 14, No. 1: JANUARY 1997, pp. 73-85

Address questions about this article to Luqi at NPS, Computer Science, Monterey, CA 93943; luqi@cs.nps.navy.mil; or to Goguen at Dept. of Computer Science and Engineering, University of California at San Diego, 9500 Gillman Drive, La Jolla, CA 92093-0114; goguen@cs.ucsd.edu.

Formal Methods: Promises and Problems

LUQI

Naval Postgraduate School

JOSEPH A. GOGUEN

University of California at San Diego

Successfully applying formal methods to software development promises to move us closer to a true engineering discipline. The authors offer suggestions for overcoming the problems that have hindered the use of formal methods thus far.

Today's fast-moving technology demands ever quicker and more reliable ways to develop software systems that meet user needs. Although industry spends billions of dollars each year developing software, many software systems fail to satisfy their users. Moreover, many systems once thought adequate no longer are, while others are never finished or never used. The September 1994 issue of *Scientific American* gives some sobering examples and concludes that "despite 50 years of progress, the software industry remains years perhaps decades short of the mature engineering discipline needed to meet the demands of an information-age society."¹ Software development failures have reached staggering proportions: an estimated \$81 billion was spent on canceled software projects in 1995 and an estimated \$100 billion in 1996.²

Many computer scientists have suggested that formal methods can play a significant role in improving this situation. Although these methods have achieved impressive successes, they have also produced disappointments. Formal methods do not yet effectively handle large and complex system development, although they can make a contribution. We know that requirements for large and complex systems are nearly always problematic initially and that they evolve continually throughout the life cycle. Thus, any method you use to implement requirements should be flexible and robust, so that it can easily accommodate the inevitable and often continuous stream of changes. We suggest that you can more effectively use formal methods by

- putting more emphasis on formal models and on domain-specific formal methods;
- using formal models as a basis for computer support of software evolution;
- using large-grain software composition methods, rather than small-grain statement-oriented programming methods; and
- taking better account of the system development context by tracing objects and relationships

back to requirements.

FORMALIZATION

Webster's Dictionary defines *formal* as definite, orderly, and methodical; defines *method* as a regular, orderly, and definite procedure; and defines *model* as a preliminary representation that serves as a plan from which the final and usually larger object is to be constructed. Thus, to be formal does not necessarily require the use of formal logic, or even mathematics. But in computer science, the phrase "formal methods" has acquired a narrower meaning, referring specifically to the use of a formal notation to represent system models during program development. An even narrower sense refers to the formalization of a method for system development. Typically, you first write a specification in a formal notation, then refine it step by step into code. Correctness of the refinement steps guarantees that the code satisfies the specification. In some methods, developers can check correctness of the refinement steps using a theorem prover for the method's underlying formal logic, but other methods remain manual because it is difficult to automate the notation used. To better understand the issues and myths related to the practical usefulness of formal methods, consult "Seven More Myths of Formal Methods,"³ and for an appraisal of their recent industrial applications, see "An International Survey of Industrial Applications of Formal Methods."⁴

Logical foundation

The prototypical formal notation is first-order logic. This notation has been extensively studied and has inference rule sets known to be sound and complete for a convenient class of models. Unfortunately, mechanical theorem provers for first-order logic can be difficult to work with.

More powerful logical systems can capture additional levels of meaning, but their theorem provers can be even harder to work with. For example, second-order logic can express security requirements for computer systems, but it does not have a sound and complete inference rule set.

Context

Experience shows that many of the most vexing problems in software development arise because any computer system is situated in a particular social context. Moreover, much of the information needed to design a system is embedded in the worlds of users and managers, and is extracted through interaction with these people. This information is informal and highly dependent on its social context for interpretation. On the other hand, we define the programming languages and other representations used to construct computer-based systems using formal syntactic and semantic rules. Both the formal, context-insensitive, and the informal, socially situated aspects of information are crucial for success. These two aspects have been called "the dry" and "the wet," and their reconciliation claimed to be the essence of requirements engineering.⁵

The dry and the wet

That we can make sense out of social life suggests it is somewhat orderly enough to be at least partly formalizable. But it is difficult to formalize domains that have many ad hoc special cases or contain much tacit knowledge or are subject to change. Formalization is more successful on narrow and orderly domains, such as sporting events, which have long traditions, regulating bodies, rule books, referees, and so on. For example, it would be more difficult to formalize a children's game than a regatta, and more difficult still to formalize human political behavior.

There are *degrees* of formalization, ranging from the very formal *dry* to the very informal *wet*. In

the driest formalizations, the metalanguage is also formalized, and an object-level model is given as a formal theory in the metalanguage. In less fully formalized models, the metalanguage may be simply a natural language, or a somewhat stylized dialect. There can be rules at both the object and the meta levels. Rules at the object level are part of the model, while rules at the meta level define the language used for formalization. For a given application, it can be a serious error to formalize more than is appropriate to the particular situation.

Formalization is useful only to the extent that it helps meet concrete goals. For example, it would only make it harder to bake cookies if the recipe were expressed in a fully formalized language. There are many similar examples in requirements engineering. Good formalizations do not usually arise top-down from desires, but rather are based on extensive experience and intuition with the domain being formalized and the intended development process.

Formal methods generally address some large class of systems, such as information systems, or even all possible systems, whereas formal models are often tailored to a specific application domain. Experience suggests that using mechanically processable formal models in building and integrating tools can yield systems that increase automation and decrease inconsistency, and thus produce software faster, cheaper, and more reliably. For example, attribute grammars are a formally processable notation that can be useful in this way. Experience also suggests using an evolutionary development process that involves rapid prototyping, such as that supported by the computer-aided prototyping system (CAPS).⁶ Such an approach contrasts with formal methods that call for mathematical rigor throughout the development process, usually by using a formal notation with a precise mathematical semantics in connection with a step-by-step refinement process. We believe it makes more sense to provide computer support for software evolution by formalizing the activities of the supporting tools rather than those of the software engineers.

In software engineering, you cannot validate results purely by proving theorems. On the contrary, you must measure the value of a contribution by its impact on practical software development and ultimately on customer satisfaction. But formalization still plays a fundamental role in software engineering, because you must have a formal (in the broad sense) model of a domain before you can design effective software for that domain. That is, problem formalization is an essential part of requirements capture.

In this respect, software engineering differs from other engineering disciplines. For example, in electrical engineering, the formalization of the problem domain is already done, and the practicing engineer need only apply it. The lack of such formalization makes software engineering more difficult than other engineering disciplines, which makes it less developed and less effectively practiced than its cousins.

Unfortunately, like many things in computer science, formal methods have been oversold. Formal methods, notations, and tools do not yet adequately support the development of large and complex systems. In general, practitioners consider formal methods useful for proving that programs satisfy certain mathematical properties, but such methods are also often considered too expensive to be practical. This view ignores evidence that appropriate and correctly used formal methods can reduce time to market, provide better documentation, improve communication, facilitate maintenance, and organize activities throughout the life cycle. Factors that influence the cost-effectiveness of formal methods include the consequences of software failure, the type of formal method to be applied, the availability of automated support for the formal method, and the skill level of available personnel.

SOFTWARE EVOLUTION

Traditionally, many in industry have viewed software evolution as occurring only after the completion of initial development. For example, L.J. Arthur defines software evolution as consisting of

"the activities required to keep a software system operational and responsive after it is accepted and placed into production."⁷ This is synonymous with maintenance, but avoids that word's negative connotations. According to Lawrence Bernstein (formerly of AT&T), evolution emphasizes the dynamic aspect of software development.⁸

Here, we consider software evolution to include all the activities that change a software system, as well as the relationships among those activities. In this case, evolution is not just another name for maintenance, because it occurs throughout the life cycle. Evolution encompasses activities ranging from adjusting requirements to updating working systems, including responses to requirements changes, improvements to performance and clarity, bug repair, version and configuration control, documentation, testing, code generation, and the overall organization of the development process.

The term "evolution" focuses attention on change. Change is inevitable and unending during software development, because it is so difficult to get a system right before it has been tried by actual users under actual operating conditions. Not only do the code and design change, the requirements and the needs that drive the requirements change as well. This occurs partly because users and analysts get a better understanding of what they really need when they see the software operating, but also because the system context changes: laws and regulations change, the competition changes, workers' expectations and habits change, management structure changes, organizational goals change, and so on.

Flexibility through prototyping

Change motivates the use of *iterative* life cycle processes, and in particular, *prototyping*: the process of quickly building and evaluating a series of concrete, executable models of selected aspects of a proposed system. In prototyping, evolution activities are interleaved with development, and continue even after delivery of the system's initial version.^{6,9}

This contrasts with traditional life cycles, such as the waterfall model, which assume that requirements can be correctly determined at the beginning of a project. Generally, project staff develop an overall system architecture from the requirements, write specifications and code for individual components, then test and debug the system. Maintenance appears in or after the final phase of testing and debugging.

Figure 1 shows an iterative prototyping life cycle. The user and designer work together to define the requirements for the envisioned system. The designer constructs a prototype at the specification level. Demonstrations of the prototype let the user evaluate the prototype's actual behavior against its expected behavior, identify problems, and work with the designer to redefine requirements. This process continues until the prototype successfully captures the critical aspects of the envisioned system. The designer then uses the validated requirements as a basis for the production software. In this way, software systems can be delivered incrementally and requirements analysis can continue throughout the system's lifetime. Incremental delivery gives users early experience with the software, leading to new goals, triggering further iterations, and extending the advantages of prototyping to the production environment.

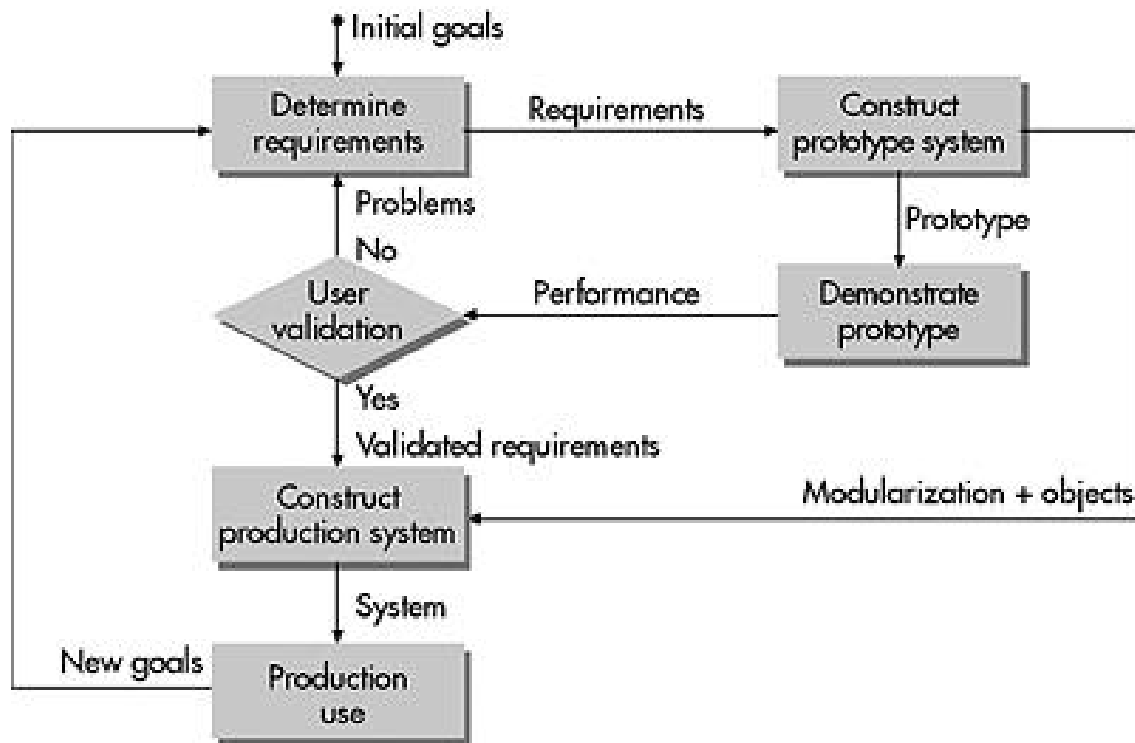


Figure 1. An iterative prototyping life cycle. Once constructed, the prototype is demonstrated to check its actual behavior against its expected behavior. This helps identify problems and can lead to a redefinition of requirements.

Evolution and formal methods

Given the inevitability of change and iteration, formal methods should be more useful in supporting evolution than in their traditional role of verifying that code meets certain fixed requirements. Possible contributions to software evolution include computer-aided design completion, program transformation, dependency maintenance (among needs, requirements, design information, documentation, code, and so on), code generation (for certain limited purposes), and merging changes to programs. These contributions become even more valuable when many programmers work concurrently on a large and complex system.

The difficulties of software evolution often extend beyond the purely technical: Social, political, and cultural factors can be significant and in many projects will dominate development costs.

Nevertheless, formal model-based tools can help maintain a software development project's integrity in many ways, such as scheduling project tasks, monitoring deadlines, tracing reasons for objects and changes, and maintaining dependency relations among versions, variations, and component decompositions. The model described in the "A Hypergraph Model for Software Evolution" sidebar is designed to support such activities.

The US Department of Defense issued MIL-STD-498 in 1994.¹⁰ This standard has evolved into ISO/IEC 12207, which will be adopted as joint standard J-STD-016. These software development

ISO/IEC 12207, which will be adopted as joint standard J-STD-016. These software development standards will have a profound effect on software evolution. They replace several previous standards that mandated the waterfall model, thus creating new opportunities by allowing considerably greater flexibility. But it also requires greater skill levels because it must be tailored to specific projects and organizations. Large and complex software development projects based on specially tailored standards will be difficult to manage without appropriate tool support, and it will be difficult to develop such tools without appropriate formal models. Contrary to typical assumptions in work on software processes, we believe that such models should focus on what the tools do rather than on what the personnel involved do.

SCALABILITY FOR APPLICATIONS

Our analysis of formal methods distinguishes between small-, large-, and huge-grain methods, referring to the size of the atomic parts used, rather than to the size of the system being developed. We reluctantly chose the word *huge* as the next step above *large*, because *large* is already in common use in certain communities; it would have been better if the three steps were instead called fine, medium, and coarse.

Small-grain methods

The classic formal methods fall into the small-grain category. These methods have a mathematical basis at the level of individual statements and small programs, but rapidly hit a complexity barrier when programs get large. In particular, systems for reasoning with pre- and postconditions such as Hoare axioms, weakest preconditions, predicate transformers, and transformational programming all have small-size atomic units and fail to scale up because they do not provide structuring or encapsulation. In general, small-grain methods have great difficulty handling changes, and thus fit poorly into the life cycle. Transformational programming is less resistant to change than other small-grain methods, but has the problem that in general there is no bound to the number of transformations that may be needed; this restricts its use to relatively small and well-understood domains.

Large-grain methods

The most important techniques of large-grain programming involve module composition. The CAPS system¹¹ provides module composition for rapid prototyping, with a dataflow-like semantics that supports hard real-time constraints and with facilities for retrieving reusable software components from a repository. The project is also working on the foundations of software maintenance and developing techniques to support design evolution, requirements tracing, configuration management, and project management. One of these techniques, change merging, has the potential to aid in combining concurrent changes to the same base version of a prototype as well as updating multiple versions of a prototype with a common change.¹²

CAPS consists of an integrated tool set that helps you design, translate, and execute prototypes. These include an evolution control system based on a graph model for evolution, a change merge facility, automatic generators for schedule and control code, and automated retrievers for reusable components.

The prototype system description language PSDL¹¹ provides a simple way to abstractly specify software systems for both prototypes and production software. A PSDL program consists of two kinds of objects, corresponding to abstract data types (PSDL types) and abstract state machines (PSDL operators) as shown in Figure 2. Their function is to localize the information for analyzing, executing, and reusing independent objects. They are also the basis for version control and are natural units of

work in a distributed implementation.

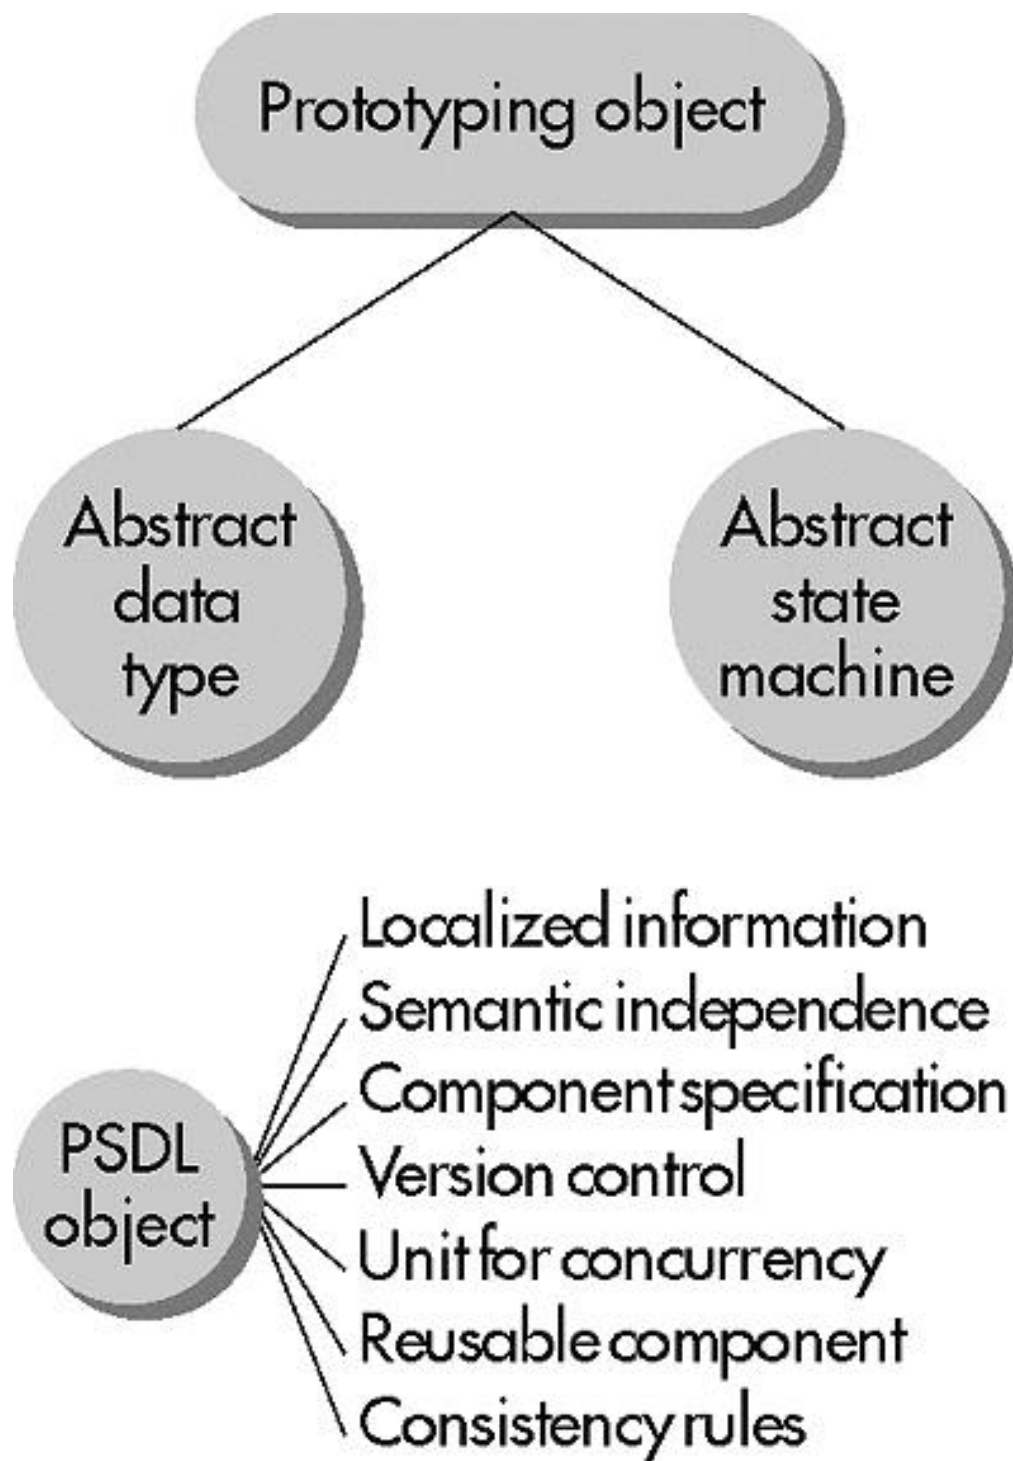


Figure 2. Class structure and properties of PSDL objects. This prototype system description language provides a simple way of specifying software systems for prototypes and production software.

When an executable Ada module is associated with each atomic PSDL object, CAPS can automatically generate "glue code" that composes these modules into a system having the structure described by the dataflow diagram. This code includes a generated schedule and tests for all the real-time constraints that have been declared; these components can be used to check the design assumptions on which the schedule is based. The system can then be compiled, executed, and tested. Error messages are produced during execution if constraints are violated.

Huge-grain methods

Huge-grain parts are much larger than small-grain statements and large-grain modules. Huge-grain parts may be systems themselves, typically commercial off-the-shelf systems. Developing systems using huge-grain parts is qualitatively different from working with small- and large-grain parts. In particular, correcting some errors in a huge-grain part may be impossible, in which case they must be accepted and worked around. For example, a network protocol such as TCP/IP may have been obtained from an external vendor, so the developers of the larger system do not have access to its source code. If the version being used has a bug, there is no choice but to find a way to avoid that bug. This is often possible because of the multiplicity of features provided in such parts. Specification and requirement methods for huge-grain systems must be robust, effective, easy to learn, and easy to incorporate into the life cycle.

Technology developed for large-grain system development can also be useful for the huge-grain case, since huge-grain parts can often be treated as modules. For example, PSDL's control constraints, such as execution guards and output guards, support adjustments to the behavior of huge-grain parts without access to their source code. The wrapper concept is also relevant. Huge-grain methods are an important area for further research.

Parameterized programming

The object-oriented version of the *parameterized programming*¹³ approach is another example of a large-grain method. It uses module expressions, theories, and views to compose systems from subsystems. It distinguishes among sorts for values, classes for objects, and modules for encapsulation. Parameterized programming lets you express designs and high-level system properties in a modular way, and lets you parameterize, compose, and reuse designs, specifications, and code as well.

In this approach, the main programming unit is the *module*, which lets you declare multiple classes together. Module composition features include summing, renaming, enhancing, modifying, parameterizing, instantiating, and importing. The *sum* of modules is a kind of parallel composition that takes account of sharing. *Renaming* lets you assign new names to the sorts, classes, attributes, and methods of modules; *enhancing* lets you add functionality to a module; and *modifying* lets you redefine some of its units.

Parameterized programming was first implemented in the OBJ language, and has also been implemented in the Functional Object-Oriented Programming System (FOOPS) and Eqlog languages. It has a rigorous semantics based on category theory. Much of the advantage of parameterized programming comes from the ability to parameterize modules using theories and views; for example, a higher-order capability can be provided in a purely first-order setting.

Parameterized programming supports *design* in the same framework as specification and coding. Designs are expressed as module expressions and can be executed if specifications that have a suitable

form are available. This gives a convenient form of prototyping. Alternatively, prototypes for the modules involved can be composed to give a system prototype by evaluating the module expression for the design. A novel feature of the approach is to distinguish between structuring, genericity, and compositionality in horizontal and vertical modes. *Vertical structure* relates to layers of abstraction, in which lower layers implement or support higher layers. *Horizontal structure* is concerned with module aggregation, enrichment, and specialization. Both kinds of structure can appear in module expressions and both are evaluated when a module expression is evaluated. The approach can also support relatively efficient prototyping through *built-in* modules, which can be composed just like other modules, and which offer a way to combine prototypes with efficient programs in a standard programming language. This is similar to the CAPS approach.

The module and type systems of parameterized programming are considerably more general than those of languages like Ada, Clu, and Modula-3, which provide only limited support for module composition. For example, interfaces in these languages can only express syntactic restrictions on actual arguments, cannot be horizontally structured, and cannot be reused. Lileanna¹⁴ implements many ideas of parameterized programming for the Ada language, including horizontal and vertical composition, following the design of the LIL (library interconnection language) system.¹³

DOMAIN-SPECIFIC FORMAL METHODS

There is much more to formal methods than suggested by the themes dominant in the past, namely synthesis and correctness proofs for algorithms. Although both of these remain interesting topics for theoretical research, their direct impact on the practice of large-scale software development is limited.

Several recent, successful applications of formal methods seem to form a cluster suggesting a new paradigm for applying formal methods. These applications involve a tool having all or most of the following attributes:

- A narrow, well-defined, and well-understood problem domain is addressed, which may have an existing, successful library of program modules.
- There is a coherent user community interested in the problem domain; the users have a good understanding of the domain, good communication among themselves, a standard terminology, and access to financial resources.
- The tool has a graphical user interface that is intuitive to the user community, embodying that community's own language and conventions.
- The tool takes a large-grain approach: rather than synthesizing procedures out of statements, it synthesizes systems out of modules; it may use a library of components and synthesize code for putting them together.
- Inside the tool is a powerful engine that encapsulates formal methods concepts and/or algorithms: it may be a theorem prover or a code generator; users do not have to know how it works, or even that it is there.

We suggest the name *domain-specific formal methods* for this emerging paradigm, in recognition of the role played by the user community and their specific domain. Some systems that fall under this heading include

- Amphion, which combines programs for astronomy calculations,¹⁵

- CAPS for real-time programming, and
- Panel for multimedia animation.¹⁶

This paradigm falls into the category of large-grain methods and can potentially be extended to huge-grain problems. The development of domain-specific formal methods should enable our discipline to replace the current practice of inventing new formal models with the more efficient practice of refining and recombining existing application models within supported domains.

This suggests a vision for the future that is less ambitious and more realistic than that of the past. It calls for using formal models and algorithms as a basis for creating computer tools to help solve practical problems that are more limited and well defined than in the past. This vision replaces the unrealistic artificial-intelligence goals of fully automatic software synthesis and verification with the recognition that human understanding and creativity must play an important role and that automated decision support can effectively enhance human capabilities. It also recognizes that requirements changes are a dominant aspect of practical software development that relies on automated tools to make software easier to change.

LIMITS AND PROBLEMS

Despite their many potential benefits, formal methods are not a panacea. We have identified nine specific problems with them¹⁷:

- Formal notation is alien to most practicing programmers, who have little training or skill in higher mathematics. Also, supporting tools are often insufficiently automated or lack user interfaces suitable for engineers.
- Formal methods papers and training often consider only toy examples taken from existing literature. Although it may be impossible to give a detailed treatment of a realistic example in a research paper or in the classroom, such examples must exist for a method to have credibility. Effective training in formal methods should treat parts of a realistic, nontrivial application.
- Many of the most popular formal methods do not scale up to practical-size problems. The gap between specifications and code is still great. Despite serious and long-term efforts in type theory, weakest preconditions, transformational programming, and so on, coding remains largely manual.
- Some advocates of formal methods dogmatically insist that everything must be proved to the highest possible degree of mathematical rigor. At the least, they argue, it must be machine-checked by a program that allows no errors or gaps, and it should be produced by a machine as well. However, mathematicians rarely achieve or even strive for such rigor; published proofs in mathematics are highly informal and often have small errors. Mathematicians never explicitly mention rules of inference from logic unless they are proving something *about* such rules. The highest levels of formality can be very expensive, and are only warranted for a system's critical aspects.
- Formal methods tend to be rigid and inflexible. In particular, it is difficult to adapt a formal proof of one statement to prove another, slightly different statement. Since requirements and specifications are constantly changing in the real world, such adaptations are frequently necessary. But classical formal methods have great difficulty in dealing with such changes; their proofs are a discontinuous function of how problems are formulated.
- Important aspects of practical software evolution are often ignored. In particular, it is

difficult to integrate formal methods into existing software processes. A related difficulty is that when you use multiple methods together you may not be able to integrate their underlying models in a way that supports building a practical software development system to support the methods.

- Often vendors do not use the best technology or even understand software development very well; they tend to be interested in profits above all and to have little time for learning either new technologies or their benefits. They often use brute-force methods to speed up projects, forgetting that this can be a shortcut to disaster.
- Some formal methods have technical difficulties. A technical deficiency of many small-grain formal methods is that first-order logic is inadequate for expressing the weakest precondition of a loop, as noted in the late 1960s by the logician Erwin Engeler.¹⁸ For example, the weakest precondition for a theorem prover for first-order logic with arithmetic cannot itself be expressed in first-order logic (the postcondition is that the input is a tautology). However, a second-order formulation is adequate, and has been used by us for some years in teaching and research, including the SPEC language used at the Naval Postgraduate School.^{9,19}
- Finally, certain fundamental limitations are imposed because all formalizations are situated in a certain context. In particular, formalizations are *emergent* in that they are always constructed and interpreted in a context. Formalizations are *contingent* in that their construction and interpretation depend upon details of the context in which this construction or interpretation actually occurs; these details may include interpretations of prior events. Moreover, interpretations are subject to negotiations among interested parties. Formalizations are *open* in that they can always be revised in the light of further analyses. They are also *vague* in that their interpretation is only elaborated to the extent that it is practically useful to do so; the rest is left as tacit knowledge. Further discussion of these points may be found elsewhere,⁵ including a general introduction to the social aspects of requirements engineering.

These limits imply that both human effort and context necessarily play a fundamental role whenever formalizations are created, interpreted, or updated. Furthermore, much of the context of that information may be social, such as goals, responsibilities, and needs associated with particular roles in an organization. These considerations are significant for designing tools to support software development. In particular, as an aid to future modifications it is highly desirable to make contextual information available along with specifications and code. The lack of such information is what makes redesign difficult and what motivates current research on reengineering. Clearly, it would be better if such information were systematically recorded in the first place.

LESSONS LEARNED

We have taken a broad view of formalization's role in the software development process and have considered the role of formal methods within that context. In particular, we found that you must understand software evolution to understand the promises and problems of formal methods, that evolution is inevitable and unending in software development, and that much of the pressure for change arises from the social context of system development. Since formal methods tend to be brittle or discontinuous—a small change in the domain can require a great deal of new work—automation is often vital for their practical application.

The construction, interpretation, and updating of formalism is always situated in a context, and understanding that context can be important in capturing the requirements for large and complex systems. Some fundamental formalization limits arise in this way. Modularity and reuse can help with any approach to improving the quality and reducing the cost of software development. We need

traceability to get better control of the software life cycle. The "Hyperrequirements" sidebar expands on this concern.

Figure 3 shows the relationships between some of the concepts we've described. A line indicates that the lower concept is a subclass of the higher one.

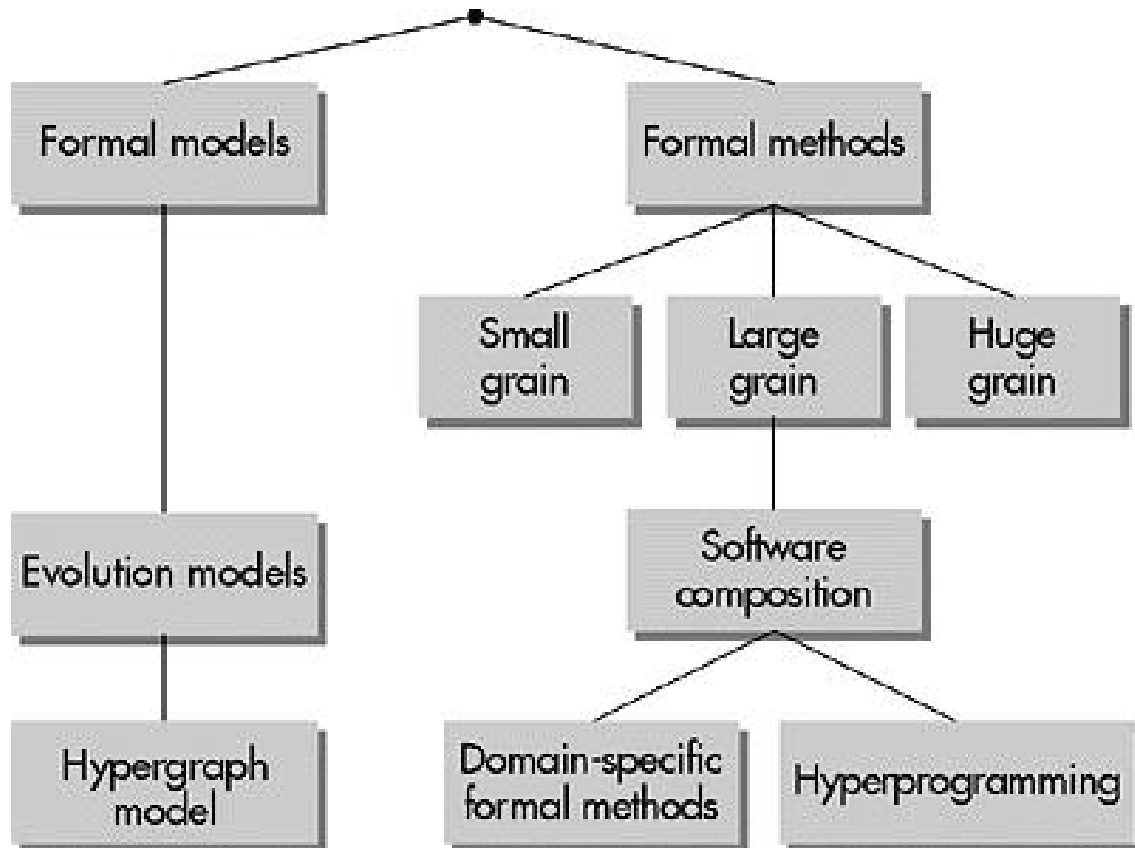


Figure 3. Classification of some formal methods concepts. Any concepts connected by a line to a higher concept is a subclass of that concept.

Building a brighter future

Whatever we learn about software development should be appropriately formalized, implemented, and put into computer science curricula so that future generations can do better than we have. Teaching a formal method while ignoring its use in real projects can have a highly negative impact. For example, students may be taught programming from formal specifications, but not that specifications come from requirements, and that requirements are always changing, often because of social, political, and cultural factors.

As a result, students are not prepared for the rapid change and political problems found in real industrial work. Many students also feel that formal methods turn programming from a creative activity into a boring, formal exercise. We have seen cases in which students have left the discipline

because teachers have failed to deal with these problems.

Students need to know how to deal with real programs that have thousands or even millions of lines of code. Most of the examples used in textbooks and the classroom are very small, however, and carefully crafted correctness proofs of simple algorithms give an entirely misleading impression of what real programming is like. Also, most of the techniques taught are small-grain and thus do not scale up to large and complex problems.

Reliable tools based on a formal model can let students do problems that would be impossible by hand. Teachers should also present methods and tools that work on large-grain units—modules rather than on small-grain units—statements, functions, and procedures—because large-grain methods can scale up, whereas small-grain methods cannot. Suites of sample problems should be developed that systematically show how and when to apply formal methods, and how to combine them with informal approaches. This will require developing appropriate module collections, refining and extending existing formal methods and tools, developing more natural user interfaces, rethinking process models, revising curricula, retraining teachers, and experimentally validating the resulting methods in practical situations.

If we fail to properly train the next generation of software developers, the problems that we see today will worsen as the size and complexity of systems continue to grow and the dead weight of legacy code continues to mount.

CONCLUSION

There is no doubt that formal models and methods can be very useful in practical software development. It also seems clear that they are necessary for transforming software engineering into a discipline that is as well understood and well organized as other engineering disciplines, which rely on sound and well-tested mathematical models. The difficulty is that formalization itself plays a more basic role in software engineering than in other engineering disciplines. Because software is still actively expanding into completely new application domains, and because requirements capture is a process of formalization, software development requires the construction of new formal models for each new application, as well as using established formal models.

More emphasis should be placed on *context* in system development and on domain-specific formal methods. However, basic research in computational logic still provides the foundation for many practical applications of formal models and methods, and advances in this area will increase the amount of computer support that can be provided in practice. A short-term view of what technology needs should be avoided, as should overselling formal methods, either as a general field or as an approach to particular applications. With these caveats, formal methods and formal models should play an increasingly important part in coming to grips with the ongoing crisis engendered by our escalating expectations about the size, complexity, and reliability of software systems.

ACKNOWLEDGMENTS

The research reported in this article has been supported in part by the National Science Foundation under grant number CCR-9058453, the Army Research Office under grant number ARO-145-91, British Telecommunications plc, the European Community under ESPRIT-2 BRA Working Group 6071, IS-CORE (Information Systems CORrectness and REusability), Fujitsu Laboratories Ltd., and a contract under the management of the Information Technology Promotion Agency (IPA), Japan, as part of the Industrial Science and Technology Frontier program "New Models for Software Architectures," sponsored by the New Energy and Industrial Technology Development Organization.

We thank Valdis Berzins and David Dampier for their valuable comments on a draft of this article. We also thank Shari Pfleeger for her inspiration and support.

REFERENCES

1. W. Gibbs, "Software's Chronic Crisis," *Scientific American*, Sep. 1994, pp. 86-95.
2. Chaos 97 tech. report, Standish Group Int'l, Dennis, Mass., to appear Jan. 1997 at <http://www.standishgroup.com/chaos.html>.
3. J. Bowen and M. Hinchley, "Seven More Myths of Formal Methods," *IEEE Software*, July 1995, pp. 34-41.
4. D. Craigen, S. Gerhart, and T. Ralston, "An International Survey of Industrial Applications of Formal Methods," TR GCR 93/626, US Nat'l Inst. of Standards and Technology, Washington, D.C., 1993.
5. J. Goguen, "Requirements Engineering as the Reconciliation of Social and Technical Issues," *Requirements Engineering: Social and Technical Issues*, M. Jirotko and J. Goguen, eds., Academic Press, London, 1994, pp. 165-200.
6. Luqi, "Software Evolution Through Rapid Prototyping," *Computer*, May 1989, pp. 13-25.
7. L.J. Arthur, *Software Evolution: The Software Maintenance Challenge*, Wiley Interscience, New York, 1988.
8. L. Bernstein, "Importance of Software Prototyping," *J. Systems Integration Special Issue: Computer-Aided Prototyping*, Vol. 6, Nos. 1-2, Mar. 1996, pp. 9-14.
9. V. Berzins and Luqi, *Software Engineering with Abstractions*, Addison-Wesley, New York, 1990.
10. Software Development and Documentation, MIL-STD-498, US Dept. of Defense, Washington, D.C., 1994, <http://www.itsi.disa.mil/cfs/std498.html>.
11. Luqi, V. Berzins, and R. Yeh, "A Prototyping Language for Real-Time Software," *IEEE Trans. Software Eng.*, Vol. 14, No. 10, 1988, pp. 1409-1423.
12. V. Berzins, *Software Merging and Slicing*, IEEE Computer Soc. Press, Los Alamitos, Calif., 1995.
13. J. Goguen, "Principles of Parameterized Programming," *Software Reusability, Volume I: Concepts and Models*, T. Biggerstaff and A. Perlis, eds., Addison-Wesley, New York, 1989, pp. 159-225.
14. W. Tracz, "Parameterized Programming in Lileanna," *Proc. 2nd Int'l Workshop Software Reuse*, IEEE Computer Soc. Press, Los Alamitos, Calif., Mar. 1993, pp. 66-78.
15. M. Stickel et al., "Deductive Composition of Astronomical Software from Subroutine Libraries," *Conf. Automated Deduction*, Vol. 12, Springer-Verlag, Heidelberg, Germany, 1994.
16. J. Schwartz and W. Snyder, "Design of Languages for Multimedia Presentations," *Proc. 1994 Monterey Workshop: Increasing Practical Impact of Formal Methods for Computer-Aided Software Development*, Naval Postgraduate School, Monterey, Calif., 1994, pp. 46-55.
17. Luqi and J. Goguen, "Some Suggestions for Progress in Software Analysis, Synthesis and Certification," *Proc. 6th Int'l Conf. Software Eng. and Knowledge Eng.*, Knowledge Systems Inst., Skokie, Ill., 1994, pp. 501-507.
18. E. Engeler, "Structure and Meaning of Elementary Programs," *Lecture Notes in Mathematics*, Vol. 188, Springer-Verlag, New York, 1971, pp. 89-101.
19. V. Berzins and Luqi, "An Introduction to the Specification Language Spec," *IEEE Software*, Mar. 1990, pp. 74-84.